

DTIC FILE COPY

(4)

198900-2-F

Final Report

HEURISTICS FOR COOPERATIVE PROBLEM SOLVING

AD-A206 371

R. FRANKLIN

L. HARMON

Sensor Systems Division

FEBRUARY 1989

Approved for Public Release
Distribution Unlimited

Prepared for:

Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

DARPA Order 5609

Contract DAAE07-87-C-R001

Issued by:
U.S. Army Tank Automotive Command
Warren, MI 48397

DTIC
SELECTED
S 21 MAR 1989 D
E



89 3 21 082

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

Form Approved
OMB No. 0704-0188

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS None										
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for Public Release Distribution is Unlimited										
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE												
4. PERFORMING ORGANIZATION REPORT NUMBER(S) 198900-2-F		5. MONITORING ORGANIZATION REPORT NUMBER(S) TR 13389										
6a. NAME OF PERFORMING ORGANIZATION Environmental Research Institute of Michigan		6b. OFFICE SYMBOL (If applicable)										
6c. ADDRESS (City, State, and ZIP Code) P. O. Box 8618 Ann Arbor, MI 48107		7a. NAME OF MONITORING ORGANIZATION U.S. Army Tank-Automotive Command										
7b. ADDRESS (City, State, and ZIP Code) Warren, MI 48397-5000												
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA		8b. OFFICE SYMBOL (If applicable)										
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Blvd. Arlington, VA 22209		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DAAE07-87-C-R001 DARPA Order 5609										
10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. 61101E		PROJECT NO.										
		TASK NO.										
		WORK UNIT ACCESSION NO.										
11. TITLE (Include Security Classification) Heuristics for Cooperative Problem Solving												
12. PERSONAL AUTHOR(S) R. Franklin, L. Harmon												
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 12/86 TO 12/87										
14. DATE OF REPORT (Year, Month, Day) 1989 February 28		15. PAGE COUNT 85										
16. SUPPLEMENTARY NOTATION												
17. COSATI CODES <table border="1"><tr><th>FIELD</th><th>GROUP</th><th>SUB-GROUP</th></tr><tr><td>9</td><td>26</td><td></td></tr><tr><td>9</td><td>2</td><td>19</td></tr></table>		FIELD	GROUP	SUB-GROUP	9	26		9	2	19	18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Cooperation, Distributed AI, Distributed Computing	
FIELD	GROUP	SUB-GROUP										
9	26											
9	2	19										
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This study examined the potential for cooperative machine intelligence. Cooperation is a common problem-solving technique in natural systems and occurs among even simple organisms. It is, therefore, a potentially powerful tool for extending the capabilities of relatively unsophisticated machines. A model of cooperation is presented which integrates problems requiring cooperative solutions, mechanisms for their implementation and the organizational requirements of cooperative groups. The biological and social science literature on cooperation is reviewed and found to yield little information about how cooperation is actually carried out in natural systems. Research in the field of distributed artificial intelligence is reviewed and critiqued. <i>This is a draft report (RF)</i>												
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified										
22a. NAME OF RESPONSIBLE INDIVIDUAL Tibor F. Czako		22b. TELEPHONE (Include Area Code) (313) 574-5515	22c. OFFICE SYMBOL AMSTA-RSS									

~~ERIM~~

PREFACE

The work described in this report represents our research on a formal model of cooperation for use in robotic systems. The program was funded by the Defense Advanced Research Projects Agency (DARPA Order No. ARPA-5609) through contract with the Department of the Army (Contract No. DAAE07-87-C-R001). We gratefully acknowledge the support and encouragement of Dr. William Isler at DARPA in pursuing this work. T. Czako at the U.S. Army Tank Automotive Command was helpful in contract administration. The findings and conclusions presented in this report do not necessarily represent the views of the above agencies.

Accession For	
NTIS GRA&I <input checked="" type="checkbox"/>	
DTIC TAB <input type="checkbox"/>	
Unannounced <input type="checkbox"/>	
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or
	Special
A-1	



TABLE OF CONTENTS

1.0	INTRODUCTION	7
2.0	BACKGROUND	9
2.1	Cooperation in the Biological and Social Sciences	9
2.2	Distributed Artificial Intelligence	12
2.3	IR&D Research at ERIM	15
3.0	A MODEL OF COOPERATION	17
3.1	Overview	17
3.2	Mechanisms for Cooperation	17
3.3	Organizational Policies in Cooperative Groups	19
4.0	CONCLUSIONS	27
5.0	RECOMMENDATIONS	29
6.0	BIBLIOGRAPHY	31
7.0	APPENDIX A: Distributed Artificial Intelligence	A-1
8.0	APPENDIX B: The Nature of Cooperation in Humans and Animals	B-1

1.0 INTRODUCTION

The work described here represents research on a formal model of cooperation, funded by DARPA Order No. ARPA-5609, Department of the Army Contract No. DAAE07-87-C-R001. Efforts were directed towards using "natural systems" for the design of robots and animated systems. Basic investigations were conducted to establish a formal model that would permit a utilitarian evaluation of information internal to a vehicle or set of vehicles from which the basis for task-oriented decision-making could be established. Planned expansion of an existing vehicle simulation to a dynamic, interactive capability which would permit testing of the formal model was not carried out when it became apparent that the dimensionality and complexity of the formal model was greater than anticipated. Limited hardware and software efforts were expended in initiating the vehicle simulation but were stopped when the magnitude of the basic investigations was understood. Since it was clear from our studies of natural systems that cooperation could substitute for individual sophistication, a formal model of cooperation became a key requirement. The major thrust of this report will address the issue of cooperation, with supporting biological and computer science studies.

To summarize briefly, cooperation is a very common problem-solving technique in natural systems and occurs in a wide variety of animals ranging from termites and ants through dogs and including humans. We humans rely on it extensively as an integral part of our societal structure. This integrated aspect of cooperation in societies (of various animals) means that research on the non-social aspect of cooperation must have careful guidelines on what to include within the scope of the research and what to exclude. However, the fact that cooperation is such a common solution means that decision-making capacities such as those of mammalian intelligences are not required for a successful implementation. Cooperation is thus a feasible area of research given the state of computing and robotics today.

A model of cooperation which included the mechanisms by which individual decision-making entities formed a cooperative group did not exist at the start of this research. Exercising this model would permit a utilitarian evaluation. An IR&D study at ERIM¹ indicated that three areas of research could be tapped to contribute to a model of cooperation of sufficient detail to suggest mechanisms for implementation: biology, social sciences including economics and policy studies, and a branch of computer science called Distributed Artificial Intelligence (DAI). Due to the nature of the studies and the style of research in the social sciences, literature which reviewed recent progress was readily available (see, for example, *The Evolution of Cooperation*²). Biology and DAI were, however, different matters. In these cases, while there was a corpus of work, it had not been systematically organized and culled, tasks which were required for the genesis of a model of cooperative behavior.

The task of reviewing relevant biological literature was delegated to Dr. Robert Taylor. His full report is included as an appendix (Appendix B). A summary of his findings is included below. We undertook an examination of the DAI literature internally, on the belief that we would be able to directly relate the elements found in the literature to the laboratory environment present at ERIM (which includes significant AI programming capabilities). The full report of the review is appended (Appendix A) and a summary is provided below.

This work has resulted in a preliminary model of cooperation which can be incorporated into decision systems of widely varying capabilities and has the potential to produce effective cooperative behaviors ranging from simple collaborations to complex, fixed hierarchies. The model is based on a concept of cooperation as a heuristic of solving problems in certain real-world situations, that is, problems of space or spatial extent, problems of time, simultaneity or concurrency, and problems of bounded rationality or logical complexity. The model of cooperative process is stated in terms sufficiently context-independent to encompass distributed computer hardware, computer software, and natural systems. Further development of the model will provide the basis for a much-needed technology of machine cooperation.

2.0 BACKGROUND

Formal studies of cooperation have occurred more or less independently in biology, the social sciences and in computer science. Biologists have investigated the conditions in which cooperation is found and its evolution as a behavior which joins individuals into groups. They have sought to explain how that cooperation could arise out of the inherently selfish process of natural selection. Social scientists have studied why and when humans cooperate, particularly from the perspective of cost/benefit tradeoffs. Most recently, computer scientists in the field of DAI have begun to consider cooperative processes as models for distributed problem-solving. Each of these disciplines provides a different view of cooperation.

No general definition of cooperation has been developed to distinguish between cooperative and non-cooperative activities of groups. Biologists, facing the enormous variety of circumstances in which cooperation arises, have used clear but limited definitions appropriate only to a particular organism and environment. Social psychologists have adopted a reward-based definition, useful for their highly structured experimental situations. For them, if a group as a whole is rewarded for a behavior, the behavior is considered to be cooperative. Within the DAI community, there is little agreement on what constitutes cooperation. Definitions of cooperation range from collective behaviors in which each member of a group selects actions on the basis of an aggregate group payoff³ to problem-solving activities which require the sharing of intermediate results among distributed problem-solvers.⁴ One of the aims of the research presented here is to develop and substantiate a general definition of cooperation which can be applied across disciplines and in different problem domains.

2.1 Cooperation in the Biological and Social Sciences

Research on cooperation in the biological and social sciences has been devoted almost entirely to questions of why cooperation occurs at all, as discussed by Taylor.⁵

Cooperation as an element of human behavior has been actively examined by social scientists for a number of years.⁶ This research has focused on why people cooperate rather than how. Even among biologists, surprisingly little attention has been given to the mechanisms by which cooperation is carried out. Instead, the major questions have been what environmental conditions give rise to altruism and cooperation among animals and why.

Human motivation for cooperation has proven to be substantially more complicated than expected. Initial game-theoretic approaches⁷ assumed that people act like economic calculators and make neutral cost-benefit analyses.⁸ This approach has produced poor predictions of human behavior. Grzelak⁹ has observed that human goals are often much more elaborate than a simple maximization of yield. For example, they may include such factors as the formation of long-term relationships or distribution of power and control within the group at the expense of short-term task performance.¹⁰⁻¹²

When humans do cooperate, they tend to form control hierarchies, to communicate extensively, and to specialize. It is important to recognize that these behaviors are often motivated by the sometimes conflicting goals of the cooperators rather than being necessary for the task at hand. Hierarchical control has its roots, at least partially, in the needs and abilities of some people to dominate and in the group's need to suppress some individually selfish behaviors. Extensive communication among humans appears to be necessary not only to accomplish a task but also to maintain cooperative attitudes.¹³⁻¹⁵ These results highlight the importance of the complex social agenda underlying and motivating human cooperative behavior. As a consequence, human cooperation may be an unnecessarily complex and misleading metaphor for machine cooperation.

Animals cooperate in much simpler ways than humans. They display two major categories of cooperation, each with many variants. Primitive cooperation is quite common among invertebrates. Social spiders build webs and attack prey together,¹⁶ enabling them to capture larger prey. They exhibit no role specialization, no obvious communication, and no coordination of behavior. Their primary manifestation of cooperation is lack of aggressiveness toward one another. Social insects exhibit significantly more integration than

spiders. Role specialization is common; an ant colony may have five different specialized castes, each of which may perform several tasks.¹⁷ They communicate with one another, usually in a primitive way with pheromones but sometimes more directly. As with social spiders, they show relatively primitive coordination of behavior.

In spite of its simplicity, invertebrate cooperation is ecologically quite successful. Its major strengths are that it requires no centralized control and that it is highly reliable. The complexities of centralized control are beyond the neural capacities of these organisms; instead, each insect has its own internal task template. The individual itself recognizes when a new job needs to be done. Reliability emerges from the simple fact that no single insect is necessary for task completion. A colony can sustain the destruction of a large fraction of its members without being eliminated as a functional unit. The major weaknesses of this kind of cooperation are that it requires redundancy in the cooperating units and does not employ any one unit particularly intelligently.

A more familiar "intelligent" form of cooperation is used by social vertebrates. Individuals cooperate to rear young, to avoid ambush, and to find and capture prey. Cooperative hunting is the most complex form of this behavior in animals. Among social hunters, role specialization is common.¹⁸⁻²⁰ Highly coordinated pursuit behaviors in pairs of falcons chasing birds suggest that integration of movements is fast and effective.²¹ Animals recognize group members individually, permitting team formation. For vertebrates, role specialization is temporary and learned rather than permanent and morphological. As a consequence, they retain more individual flexibility.

While animal cooperation may turn out to be as complex and confusing as human cooperation, this seems unlikely. Many of the social problems that concern people are unimportant to, or at least radically different for, animals in groups. Whether animals provide useful models for machine cooperation requires a good deal of thought; still, it is clear that the behaviors of cooperating humans are not necessary for simple tasks. Those of cooperating wolves or wasps might perform more reliably and could be realized at a fraction of the cost. As is discussed below (Section 2.3), natural systems are sources of examples rather than strict models for imitation.

2.2 Distributed Artificial Intelligence

The goal of DAI research has been to implement general problem-solving heuristics for cooperative interactions among computers.²² DAI is concerned with the use of logically distributed components to solve problems which are perceived to be too large for a single component but which do not decompose readily in a fixed form.²³ Problem-solving components are variously termed agents, nodes or knowledge sources, depending upon the author's perspective. Intuitive notions about how groups of human experts solve complex problems form the standard DAI model.²⁴ Central issues have been problem decomposition, relationships among problem-solving nodes (communication and control systems), and single-agent reasoning about multiple agents. Although a variety of implementation approaches and much detail may be found in the DAI literature, no organizing principles have emerged. DAI research has proceeded by applying arbitrary models of cooperative processes to specific problems.²⁴⁻²⁶ As initial models and assumptions have broken down, incremental, compensatory improvements have been attempted, with considerable effort. This approach has limited the generality of results so that solutions to closely related problems may have to be found independently.

Three representative DAI simulation systems illustrate the major issues addressed in the DAI literature. The Distributed Vehicle Monitoring Testbed (DVMT) of Lesser, Corkill, and co-workers^{4,27-30} is a simulated distributed sensing network designed for empirical studies of control organization and communication policies. In the DVMT, sensors are distributed over an area containing genuine and spurious vehicle tracks which are stationary but extend beyond the sensing range of any single sensor. The network task is to arrive at descriptions of genuine tracks, requiring the exchange and integration of partial results from different sensors. The Distributed Sensing System (DSS) of Smith and Davis^{24,31} is a simulated network of stationary specialist nodes; each can sense, process sensory data or integrate results. The network task is to identify vehicles from simulated acoustic signals and track them during vehicle motion. The DSS is designed to study the Contract Net, a negotiation protocol for matching agents with sub-tasks and assembling subproblem solutions into global solutions. The Contract Net was modeled after groups of

human experts and is intended for problems which decompose naturally into hierarchically organized subproblems. RAND's ongoing program of Remotely Piloted Vehicle (RPV) simulations^{26,32} grew out of their work in distributed Air Traffic Control.³²⁻³⁵ For the RPV simulations, RAND added simple surveillance missions for groups of autonomous aircraft in benign and hostile environments.

Both the DVMT and DSS address problems of distributed interpretation²⁶ rather than distributed action. In each case, interaction among agents occurs strictly through communication. Although the DSS involves the interpretation of sensor input from a dynamic environment, nodes do not alter the environment. The RAND work differs in that agents not only observe the environment but move within and alter it; as a result, agents interact not only via communication but also through their effects on the problem environment. In this respect, RAND has moved closer to investigating the requirements of cooperative agents in the real world.

Coherence is a theme which pervades the DAI literature: how can the activities of separate agents be integrated into a single problem solution? "Coherence," as it is used in DAI, closely parallels the idea of means-end coherence,³⁶ the requirement that a single agent plan be reasonably expected to lead to the desired end. However, the relationship between the coherence issues in single agent and multiagent problem-solving remains unclear. Based upon a paradigm of human experts as problem-solvers, coherence was attempted through centralized control supported by extensive communication. The underlying assumption was that increasing centralization reduced communication requirements. Research proceeded by selecting particular (primarily centralized) organizational structures, adding flexibility in response to empirically demonstrated needs, and tailoring communication policies to minimize overhead in the resulting systems.

The RAND RPV simulations started with a fixed hierarchical organization,³² but the effectiveness of a fixed hierarchy was found to drop with increasing environmental complexity. Organizational flexibility through role reassignment was added later.²⁶ The DSS of Davis and Smith is a flexible hierarchy based on negotiation. Davis and Smith concluded that the communication and computation overhead of a flexible hierarchy is

justified only when subproblems are large, and when natural hierarchies of task or of data abstraction exist. A formally decentralized control structure was employed in most of the DVMT experiments, although the effective organizational structure was varied. No single organizational structure was found which allowed the DVMT network to function optimally under all conditions.

In accord with the emphasis on centralized organizational structures, cooperative planning in the DAI literature has mainly been centralized, i.e. one agent planning the actions of many. The hope was that centralized multiagent planning systems could be developed based on single agent planners. Issues associated with centralized planning include the ability to represent concurrent events, to identify spatial and temporal relationships among concurrent events, especially potential interference between the actions of multiple agents, and to represent and reason about the actions of other agents. Centralized planning, however, implies reasoning about the actions of agents over which the planner has control. Therefore, solutions to the centralized multiagent planning problem will not necessarily provide solutions for genuinely distributed planning.

Researchers in the field agree that robust and coherent distributed problem-solving requires more flexibility than is possible within current approaches. As a way to introduce flexibility, reduce overhead and increase coherence, DAI researchers have expanded the kinds of knowledge shared among agents. A basic level of shared knowledge was incorporated in DAI research at the outset and termed the benevolence assumption. This is the assumption that all agents are working toward the same goal. The need for additional shared knowledge has been demonstrated empirically in each of the major DAI simulations. The emphasis in the literature has been on the sharing of information about the problem domain, of what agents *know*, and knowledge about agent capabilities (what other agents can or will *do*). Researchers appear to be looking for agent models which will allow agents to predict and account for the actions of other agents during problem-solving. No systematic study of agent models and heuristics for their use in problem-solving has yet been published.

A bottom-up approach, beginning with particular problems and applying specific (human) models, has characterized DAI research. With human problem-solving as a paradigm, hierarchies were an obvious choice for group organization, although other forms of cooperation exist which solve problems without the expense of hierarchies. Discrete issues are being addressed, such as the value of negotiation protocols and mechanisms for implementing them, but no organizing principles have emerged from the work done to date. The current state of the field suggests that a higher level view of distributed problem-solving will be necessary in order to make significant advances in the field and to consolidate gains made thus far.

2.3 IR&D Research at ERIM

Research at ERIM has investigated cooperation as a fundamental problem-solving heuristic, for the purpose of identifying underlying principles. Such principles will form the foundation for future realization of cooperation in machines. The ERIM study started with animal rather than human cooperative systems. The study of natural systems is not intended to lead to machine replication or reverse engineering. Rather, natural systems are examples of demonstrably successful solutions to real-world problems. From the rich array of alternatives in nature it has been possible to extract general features of cooperative problem-solving, as well as the specifics of their application in real-world domains. Because all human cooperative behaviors have pronounced social aspects, it is impossible to separate social and non-social functions in human models. The variety in natural systems makes comparative studies possible. The use of animal models should therefore simplify the identification of non-social requirements of problems solved by cooperation.

A functional definition of cooperation may be developed from its effect on group performance and the problems which it solves. Our preliminary model of cooperation derived from natural systems suggests that cooperation is a set of heuristics used by a group of agents in order to extend their abilities in a synergistic fashion. Cooperation overcomes specific types of individual limitations. The first of these is spatial; individuals can only operate over limited areas. Individuals are also limited in the number of actions they can

perform simultaneously. The final limitation lies in the finite number of different skills any individual can apply to a problem. This limitation is also referred to as "bounded rationality," a term originating from studies of human economic behavior which refers to problems which require more knowledge or analytical power than a single agent can provide. This leads to the definition of three problem classes for which cooperation is appropriate:

- problems of spatial extent,
- problems of simultaneous action, and
- problems of bounded rationality.

This research has resulted in a preliminary model of cooperation which integrates problem classes, mechanisms and group characteristics into a single theory.¹

3.0 A MODEL OF COOPERATION

3.1 Overview

Combining the result from the preliminary study with the results of our background investigations has led us to propose a robust model of cooperative behavior. The basic elements of this model are presented in Figure 1. First, agents require skills to solve specific tasks in the problem domain, referred to as algorithms. In order to exercise those skills cooperatively, agents need to be able to share information via communication and/or inference, and to coordinate their actions in space and time. These mechanisms are discussed in more detail in Section 3.2. Heuristics, knowledge about the use of these mechanisms and of particular agent skills, are also required for cooperation. At the highest level, protocols must be shared concerning how decision-making is to be carried out within the group, which constitutes the group's organizational policy. Organizational policies are closely related to the distribution of algorithms and heuristics within groups. Therefore these elements of the model are discussed together in Section 3.3. While this model appears to be self-consistent and our work has shown its explanatory power, considerable research is needed to give it the power and robustness desired in a comprehensive theory of cooperation.

3.2 Mechanisms for Cooperation

A set of components (mechanisms) has been defined which appear to be necessary and sufficient for cooperative action. These definitions were tested through analysis of existing examples of cooperation and by implementation in a computer simulation of cooperative hunting. The basic cooperative mechanisms are postulated to be coordination, communication and inference.

Coordination is the concerted application of individual skills by multiple agents during problem solution. It entails the intentional correlation of actions in time and space. Coordination may arise in diverse ways. For example, it may come about through execution

Elements of a Model of Cooperation

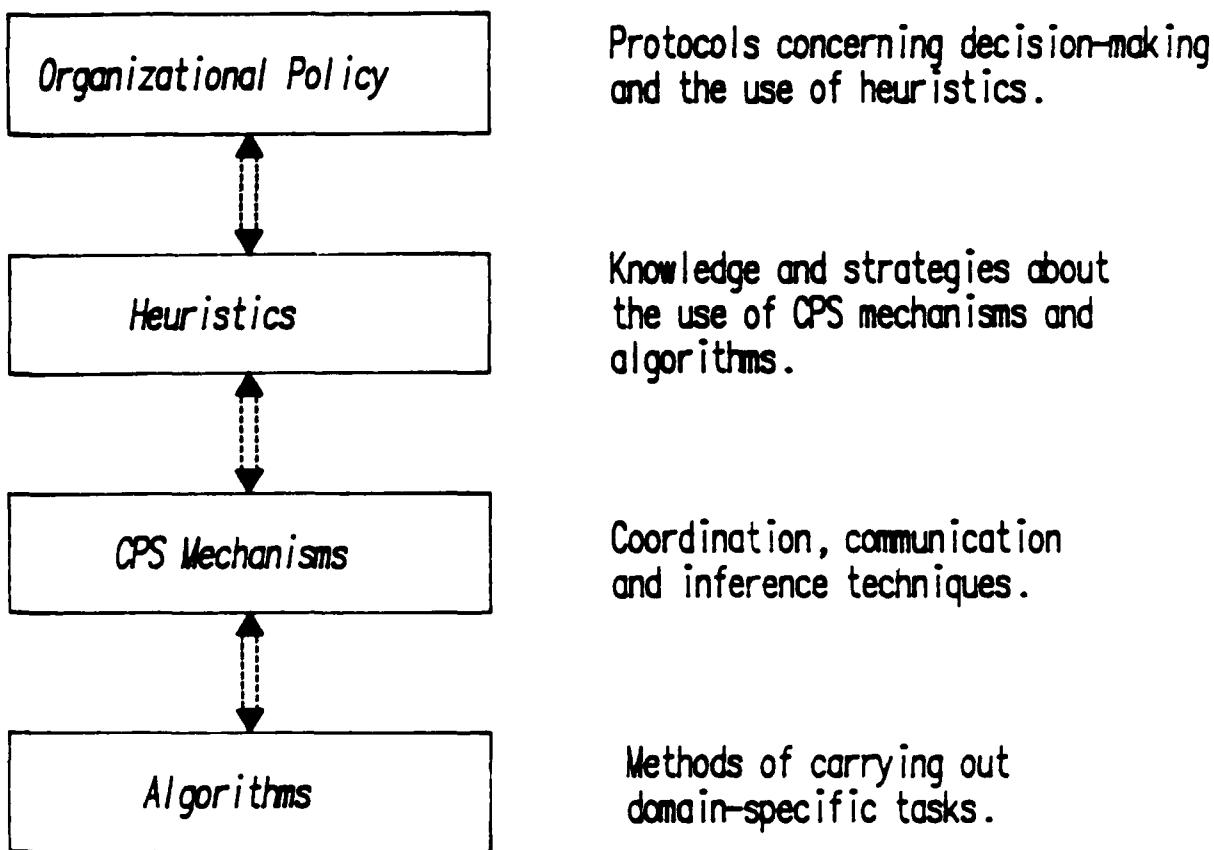


FIGURE 1.
Elements of a Model of Cooperation

of a detailed common plan or may unfold dynamically strictly in response to the problem environment and other agents. In either case, coordination is the means by which the actions of agents are combined to implement a particular cooperative strategy. While recognizing that intentionality is a complex issue,³⁷ "intention" is used to exclude from coordination those actions whose apparent correlation is due only to chance. An important aspect of coordination is that coordinated actions must occur within some spatial or temporal interval of each other. The maximum allowable spatio-temporal separation of coordinated acts establishes a coordination window--determined by agent properties, such as memory degradation over time or limitations in communication and sensing ranges, as well as by the requirements of a specific task. From the perspective of the observer, however, an apparent coordination window may be limited by the observer's ability to perceive correlations.

Information transfer may occur by communication or inference alone, or by some combination. For the purposes of analyzing cooperative interactions, it is useful to define communication and inference as mutually exclusive processes. Communication is the passage of a message from a transmitter to a receiver and is determined by the intent of the sender without regard to the receiver status. This definition is valuable when considering action-oriented processes in spatially and temporally complex domains which leave the sender uncertain as to the status of potential receivers. Inference, on the other hand, is defined as the process whereby one agent obtains information through observation of another in the absence of communication. It is therefore a receiver-based process. Inference requires two things of the receiver: first, it must have a way to monitor the actions of another agent; and second, the observer must have some knowledge of the rules by which the observed agent is generating its behavior. This type of inference is also referred to in the literature as non-communicative plan recognition.

3.3 Organizational Policies in Cooperative Groups

Apart from the mechanisms which support cooperative action, described above, the remaining elements of the model are properties of groups. Three essential characteristics

of cooperative groups can be identified by examination of a variety of cooperative systems: 1) the distribution of domain-specific skills; 2) the distribution of knowledge about the effects and use of those skills in problem-solving; and 3) the degree to which decision-making is centralized. Domain-specific problem-solving skills we refer to as algorithmic knowledge--knowledge of distinct methods to carry out particular tasks. Knowledge about agents and the effects of their skills on problems we refer to as heuristic. Decision-making in our terminology is synonymous with control--the ability to determine what actions will be taken, or which algorithms will be used and when. The way decision-making is distributed defines the "organizational policy" of a group.

These characteristics can be organized into a three-dimensional graphic model, in which each is represented on a separate axis (Figure 2). On one axis, the distribution of algorithms is displayed. At one end of this axis, all group members share the same algorithms, a homogeneous distribution of skills. At the other extreme, agents are highly specialized and different from one another, a heterogeneous distribution. The distribution of heuristic knowledge is represented on a second axis. Heuristics may be shared equally within a group or members may have distinctly different knowledge, so that the distribution of heuristic knowledge also ranges from homogeneous to heterogeneous.

Decision-making, shown on the vertical axis, may be fully decentralized (or distributed) or fully centralized, with considerable variation between the two limits. The distribution of heuristics determines the minimum possible centralization of decision-making: decision-making must be restricted at least to those agents with the requisite heuristic knowledge. This dependence is indicated by the curve drawn in the Decision-Making/Heuristics plane in Figure 2. This diagram is intended to suggest qualitative relationships rather than precise functional dependencies.

Specific organizational policies found in natural systems can be located in this graphic representation, as shown in Figure 3. For example, *heterarchies* are groups of agents with identical problem-solving skills, so that algorithms are homogeneously distributed. Each agent can use itself as a model in predicting and interpreting the actions of others. Therefore, problem-solving heuristics are homogeneously distributed as well. No gain is

Organizational Policies

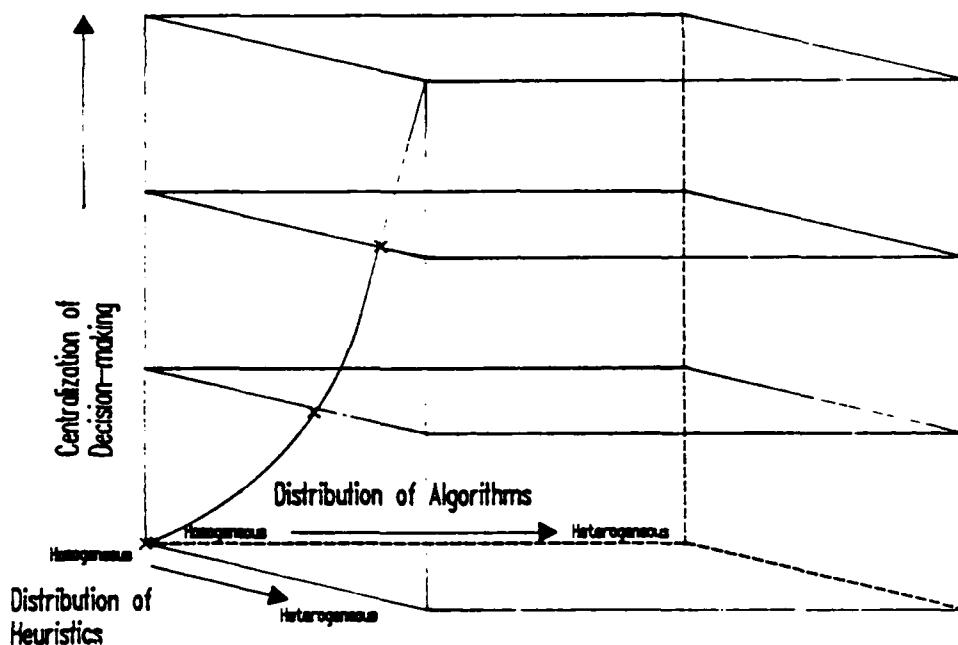


FIGURE 2.
Organizational Policies in Cooperative Groups.

Three essential characteristics define the space of possible cooperative groups: 1) the distribution of domain-specific skills (algorithms); 2) the distribution of knowledge about the effects and use of those skills in problem-solving (heuristics); and 3) the degree to which decision-making (control) is centralized. The way control is distributed determines the "organizational policy" of a group.

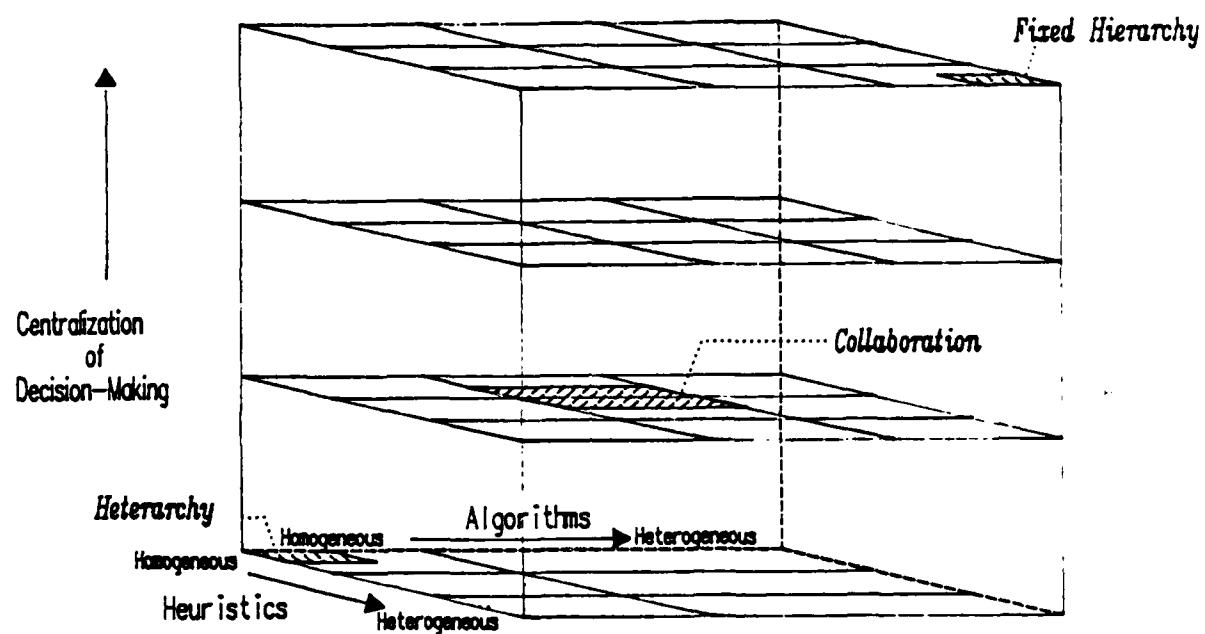


FIGURE 3.
Sample Organizational Policies.

obtained by having a single individual among a group of identical agents plan for the entire group. Heterarchies therefore operate without centralized control, placing them at the origin of the diagram in Figure 3.

The term *collaboration* is used to describe groups which are somewhat heterogeneous in skills, and in which there is some heterogeneity in the distribution of heuristic knowledge, leading to a small degree of centralization in decision-making. This places them near the middle of Figure 3. When the heuristics necessary for decision-making are localized in a few agents, for example planning and integration specialists, control is necessarily centralized and a *fixed hierarchy* results, as indicated on the upper right.

In Figure 4, additional cooperative forms are added to those of Figure 3. A *coalition* is a loose association of somewhat specialized agents and requires extensive overlap in heuristic knowledge for agents to integrate their own activities during problem solution. For social insects, extreme skill specialization is associated with a natural and complete problem decomposition. For example, foraging and defense are independent activities. What is shared among skill specialists is the ability to recognize the need for their skills. "Subproblems" of the overall problem of group survival are integrated not through a central decision-maker but through the environment, through the effects of each on the problem. This permits a group of skill specialists to operate in a coalition.

Flexible hierarchies are typically composed of agents specialized both in skills and heuristics. The assignment of control in a flexible hierarchy must be transferable among agents; therefore, although a flexible hierarchy entails centralization of decision-making, the necessary heuristics must be shared at least to some extent, which places them to the left of fixed hierarchies in Figure 4. Agents which are extremely specialized in problem-solving heuristics require centralized control because either: 1) all requisite knowledge is localized in one or a few agents, or 2) the results of different heuristics distributed among the agents must be integrated by another. Therefore, the region of Figure 4 in which heuristics are heterogeneously distributed but control is *de-centralized* is labeled *anarchy*; cooperation will not be effective in such a system.

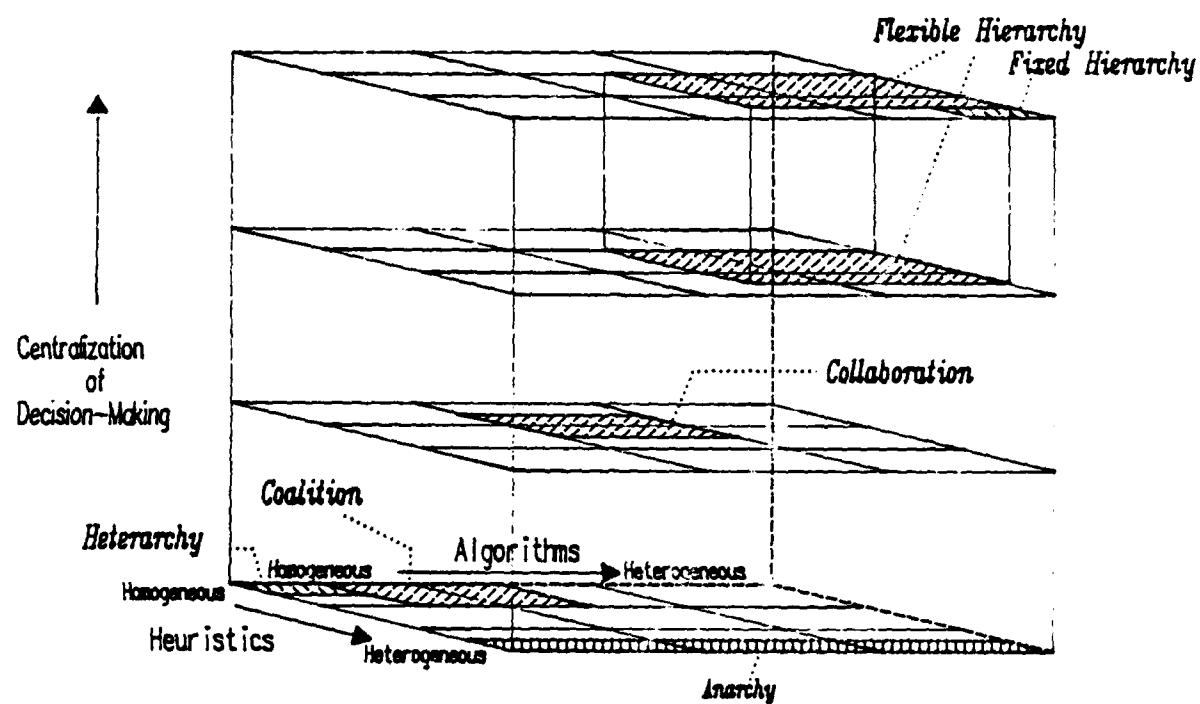


FIGURE 4. Organizational Policies in Natural Systems.

Another term for groups which are heterogeneous, in either domain-specific skills or heuristics, is "specialized". The degree of specialization in those two aspects is critical in determining organizational policies. While the word "specialization" commonly refers to inherent properties of individuals or *intrinsic specialization* (such as mobility characteristics or sensing capabilities), another kind of specialization exists. The unique knowledge and capabilities acquired by an individual agent during problem solution give rise to *extrinsic specialization*, specialization induced by the interactions of agents with the problem environment. For example, identical vehicles conducting a search will gain knowledge about different parts of the terrain, thus becoming *extrinsic* terrain specialists. To our knowledge, the distinction between intrinsic and extrinsic specialization has not been addressed elsewhere, although both affect the choice of appropriate group structures and the necessity for centralized decision-making.

From the model described here, the individual requirements of cooperative agents can be determined. The model offers a framework for integrating: 1) problems requiring cooperative solutions; 2) mechanisms for their implementation; and 3) the organizational requirements of cooperative groups. This satisfies a clear need for a unified theory of cooperation. The preliminary theory therefore provides the basis for a technology of cooperation detailing how cooperative machine intelligence should be implemented and employed in solutions to specific problems, based on problem characteristics, available agents and an assessment of costs and benefits.

4.0 CONCLUSIONS

The research reported here has shown that cooperation is a fundamental procedure-domain heuristic. It is clear from the reviews conducted in this program that no basic principles exist to guide the development of cooperative machines. In order to fill that gap, we have now established the beginnings of a formal theory and descriptive model of cooperative problem-solving. When completed, this model will establish the required capabilities of cooperative agents. It will provide a basis for a technology of cooperation detailing how cooperative machine intelligence should be implemented and employed in solutions to specific problems, based on problem characteristics, available agents and an assessment of costs and benefits.

5.0 RECOMMENDATIONS

Continued development of a technology of cooperation will require two parallel and mutually supportive lines of research: concept development and computer simulation. Principles derived from natural systems will be a primary input for concept development. As concepts are developed, they should be implemented and tested through computer simulations of specific cooperative missions in complex and dynamic environments. These simulations should be evaluated from a cost-benefit standpoint and cooperative performance compared with that attained by solitary agents.

A critical question to be addressed is how much autonomy individual agents require in order to cooperate. A set of domain-independent mechanisms has been postulated to be necessary and sufficient for cooperative behavior. These mechanisms require further definition and understanding before implementation in cooperative machine interactions. The different forms of cooperation must be clearly distinguished from each other and from non-cooperative activities of groups. Finally, the forms of cooperation should be ranked according to ease of implementation and the relationship between complex and simple forms examined.

The effectiveness of the various forms of cooperation on different types of problems must be investigated. Problems of spatial extent and simultaneous action should be investigated early. The results of our research indicate that cooperation can be successful in these problem classes with minimal individual sophistication. Research is needed to investigate principles of cooperative problem-solving for heterogeneous groups, those comprised of agents with only partial overlap of skills or knowledge. A firm theoretical understanding of the relationship between cooperation and group heterogeneity must be established. Following this, problems of bounded rationality, which require specialists, can be studied in a systematic fashion.

The research recommended here would serve to test the limits of the initial findings presented in this report and would lay the groundwork for the use of cooperation in a wide variety of practical systems.

6.0 BIBLIOGRAPHY

1. R. F. Franklin and L. A. Harmon, Elements of Cooperative Behavior, ERIM Report No. 655404-1-F (1987).
2. R. Axelrod, The Evolution of Cooperation. Basic Books, New York, NY (1984).
3. M. R. Genesereth, M. L. Ginsberg, and J. S. Rosenschein, Cooperation without communication, AAAI Conf. Proc. 1, 51(1986).
4. V. R. Lesser and D. D. Corkill, Functionally Accurate, Cooperative Distributed Systems, IEEE Trans. on Sys., Man. and Cybern. SMC-11, 81(1981).
5. R. J. Taylor, The nature of cooperation in humans and animals: A report for the Environmental Research Institute of Michigan, 1987.
6. V. J. Derlega and J. L. Grzelak, Cooperation and helping behaviors: theories and research. Academic Press, NY (1982).
7. A. M. Colman, Cooperation and Competition in Humans and Animals. Van Nostrand Rheinhold, Berks (1982).
8. M. Deutsch, Interdependence and psychological orientation, in: V. J. Derlaga and J. L. Grzelak, Cooperation and helping behaviors: theories and research, pp. 15-42.
9. J. L. Grzelak, "Preference and cognitive process in interdependent situations: a theoretical analysis of cooperation"; in: V. J. Derlaga and J. L. Grzelak, Cooperation and helping behaviors: theories and research, pp. 95-122.
10. O. Ben-Yoav and D. G. Pruitt, Resistance to yielding and the expectation of cooperative future interaction in negotiation, J. Exper. Soc. Psych. 20, 323(1984).
11. I. L. Janis, Groupthink: psychological studies of policy decisions and fiascoes. Houghton Mifflin, Boston (1982).
12. B. Latane, S. A. Nida and D. W. Wilson, "The effects of group size on helping behavior", in: J. P. Rushton and R. M. Sorrentino, eds., Altruism and Helping Behavior. Erlbaum Assoc., NJ (1981).
13. S. L. Braver and L. A. Wilson II, Choices in social dilemmas: effects of communication within subgroups, J. Confl. Resol. 30, 51(1986).

14. W. B. G. Liebrand, "The effect of social motives, communication and group size on behaviour in an N-person multi-stage mixed-motive game", Eur. J. Soc. Psych. 14, 239-264 (1984).
15. D. M. Stults and L. A. Meese, Behavioral consistency: the impact of public versus private statements of intentions, J. Soc. Psych. 125, 277-278 (1985).
16. R. Buskirk, Sociality in arachnida. In H. Herman, (ed.) Social Insects, pp. 282-369. Academic Press, NY (1981).
17. E. O. Wilson, The insect societies. Harvard Univ. Press, Cambridge, MA (1971).
18. K. L. Hamlin, Cooperation by coyote pairs attacking mule deer fawns. J. Mammal. 60, 849-850(1979).
19. D. L. Mech, The wolf: the ecology and behavior of an endangered species. Univ. Minn. Press, Minneapolis, Minn. (1970).
20. G. B. Schaller, The Serengeti Lion: a study of predator-prey relationships. Univ. Chicago Press, Chicago, Ill. (1972).
21. D. P. Hector, "Cooperative hunting and its relationship to foraging success and prey size in an avian predator", Ethol. 73, 247-257 (1986).
22. L. A. Harmon and R. F. Franklin, "Distributed artificial intelligence: a literature review and critique", SPIE Vol. 1008 Expert Robots for Industrial Use, 136 (1988).
23. R. Davis, Report on the Workshop on Distributed AI, IEEE SIGART Newsletter 73, 42 (1983).
24. R. Davis and R. G. Smith, Negotiation as a Metaphor for Distributed Problem Solving, Artificial Intelligence 20, 63 (1983).
25. M. N. Huhns, L. M. Stephens, and R. D. Bonnell, Control and cooperation in distributed expert systems, IEEE Southeascon, 241(1983).
26. R. Steeb, D. McArthur, S. Cammarata, S. Narain, J. Rothenberg and W. Giarla, Cooperative intelligence for RPV fleet control: analysis and simulation, The Rand Corporation, Working Draft, WD-2598-ARPA, (April, 1985).
27. D. D. Corkill and V. R. Lesser, The use of meta-level control for coordination in a distributed problem solving network, IJCAI 83, 748 (1983).
28. E.H. Durfee, D. D. Corkill and V. R. Lesser, Distributing a distributed problem solving network simulator, Proc. Real-Time Systems Symposium 1984, 237(1984).

29. E. H. Durfee, V. R. Lesser and D. D. Corkill, Coherent cooperation among communicating problem solvers, Draft, (September, 1985).
30. V. Lesser, D. Corkill, J. Pavlin, L. Lefkowitz, E. Hudicka, R. Brooks, and S. Reed, A high-level simulation testbed for cooperative distributed problem solving, IEEE 3rd Int. Conf. on Distributed Computing Systems, 341(1982).
31. R. G. Smith and R. Davis, Frameworks for cooperation in distributed problem-solving, IEEE Tran. on Sys., Man and Cybern. SMC-11, 61 (1981).
32. R. Steeb, D. McArthur, S. Cammarata, S. Narain, and W. Giarla, Distributed problem solving for air fleet control: framework and implementation, The Rand Corporation N-2139-ARPA, (April, 1984).
33. S. Cammarata, D. McArthur and R. Steeb, Strategies of cooperation in distributed problem solving, The Rand Corporation N-2031-ARPA, (1983).
34. D. McArthur, R. Steeb and S. Cammarata, A framework for distributed problem solving, Proc. Nat. Conf. Art. Intell., 181 (1982).
35. R. Steeb, S. Cammarata, and S. Narain, Distributed intelligence for RPV fleet control, Proc. IEEE EASCON, 223 (1984).
36. M. Bratman, Taking plans seriously, Social Theory and Practise 9, 271(1983).
37. P. R. Cohen and H. J. Levesque, Persistence, Intention, and Commitment, SRI Tech. Note 415, (1987).

APPENDIX A

Distributed Artificial Intelligence: A Literature Review and Critique.

SUMMARY.

The goal of DAI research has been the development of general problem-solving heuristics for cooperative interactions among computers. It is one of two major lines of research concerned with distributed, multi-computer solutions to complex problems, the other being distributed processing (DP). DAI is concerned more with the use of logically distributed components or agents to solve problems. These are problems which are perceived to be too large for solution by a single component but which are not readily decomposable in a fixed form [2]. DAI research has been focused on such issues as relationships among problem-solvers, and the planning of multiple agent activities by a single agent.

Although a variety of implementation approaches and wealth of detail may be found in the DAI literature, no organizing principles have emerged to date. We attribute this lack to the way the research has been carried out, implicitly or explicitly applying specific (human) models of cooperative process to particular, practical problems. This approach carries the risk that results will be of limited applicability and solutions to similar problems will have to be found de novo. As models and initial assumptions have broken down, considerable effort has been expended on incremental, compensatory improvements to individual cooperative policies. The state of the field today suggests that a higher level view of distributed problem-solving will be necessary in order to make significant advances in the field and to consolidate gains made thus far.

In this review, three representative DAI simulation systems are described: 1) the Distributed Vehicle Monitoring Testbed (DVMT) of Lesser, Corkill, and co-workers [3,4,5,6,7]; 2) the Distributed Sensing System (DSS) of Smith and Davis [8,9]; and 3) the RAND Remotely Piloted Vehicle (RPV) simulations [10,11]. These simulations illustrate the major issues addressed in the DAI literature.

Coherence is a theme which pervades the DAI literature, i.e. how can the activities of separate agents be integrated into a single problem solution? Coherence as it is used in DAI closely parallels the idea of means-end coherence [16], the requirement that elements of a single agent plan be reasonably expected to lead to the desired end. However, in DAI research it has been phrased strictly as the problem of obtaining globally relevant action from multiple agents having only local knowledge [2,4]. Little effort has been devoted to analysis of the relationship between the coherence issues in single agent and multi-agent problem-solving. Existing DAI

systems have adopted specific *ad hoc* structures which originate from human models in particular domains [9,10,19]. Based upon a paradigm of human experts as problem-solvers, initial attempts to ensure coherence have relied on centralized control, supported by extensive communication. As a consequence, the assumption has been made that hierarchical structures will, in general, require less communication than less centralized control structures. An overall theory is lacking by which to determine the veracity of this assumption by testing the relative control requirements and corresponding communication needs of distributed solutions to different types of problems. Research has proceeded by selecting particular (primarily centralized) organizational structures, adding flexibility in response to empirically demonstrated needs, and tailoring communication policies to minimize overhead in the resulting systems.

Coherence leads directly to considering forms of organization. The RAND RPV simulations started with a fixed hierarchical organization [11], but their results indicated that optimal group performance could then be obtained only in simple environments. Default lines of succession, assignment by group leader and negotiation were added to introduce flexibility by role reassignment [10]. In general, the effectiveness of a fixed hierarchy was found to drop with increasing environmental complexity. The DSS of Davis and Smith is organized as a flexible hierarchy based on negotiation (Davis and Smith have pointed out that negotiation is an expensive protocol). They concluded that the communication and computation overhead involved in a flexible hierarchy is justified only when subproblems are large, and natural hierarchies of task or of data abstractions exist. A formally decentralized control structure was employed in most of the DVMT experiments, although implicit organizational structure was varied. However, of the forms tested, no single organizational structure was found which allowed the DVMT network to function optimally under all conditions.

Communication has been used to enhance coherence by increasing the knowledge of individual agents, both in the problem domain and about each other. An important practical question has been how to minimize the resulting communication loads. Communication reduction has been attempted by: limiting message circulation (the number of recipients) and limiting message content (the type of information to be communicated). The two extremes of a circulation mode are broadcast and directed transmission. Of the two, broadcast is the most commonly implemented form in the DAI literature. Durfee, *et al.* [7], have demonstrated a complex interaction between the type of information transmitted, the weight given to received information by the receiver, and specific problem characteristics in determining system performance.

In accord with the emphasis on centralized organizational structures, cooperative planning in the DAI literature has primarily referred to centralized planning, that is, planning by one agent for the actions of multiple agents. The hope has been that centralized multi-agent planning systems could be developed using as a basis previous work on single agent planners. Elements which have been associated with centralized planning include the ability to represent concurrent events, to identify spatial and temporal relationships among concurrent events, especially potential interference between the actions of multiple agents, and to represent and reason about the actions of other agents. Centralized planning, however, implies reasoning about the actions of agents *over which the planner has control*. Therefore, solutions to the centralized multi-agent planning problem will not necessarily provide solutions for distributed planning.

Researchers in the field agree that robust and coherent distributed problem-solving will require more flexibility than is possible within current approaches. As a way to introduce flexibility, reduce overhead and increase coherence, DAI researchers have expanded the kinds of knowledge shared among agents. A basic level of knowledge about distributed problem-solving implicitly shared among agents was incorporated in DAI research at the outset, the benevolence assumption, i.e. the assumption that agent actions are directed to the same goal. The need for additional shared knowledge, has been demonstrated empirically in each of the DAI simulations. The emphasis in the literature has been primarily on the sharing of information about the problem domain, of what agents *know*, and knowledge about agent capabilities (what other agents can or will *do*). Researchers appear to be looking for agent models which will allow agents to predict and account for the actions of other agents during problem-solving. No systematic study of agent models and heuristics for their use in problem-solving has yet been published.

Introduction.

A literature review normally presents the published material of a field within an organizing framework, either in some explicit form such as that of a major theory, or implicitly as suggested by the outstanding problems in the body of work. Research in the field of Distributed Artificial Intelligence (DAI)¹ does not readily admit to this form of organization. It is fragmented, with few guiding principles or priorities. The fragmentation arises from the way in which the foundations for the research have been established. By relying on a bottom-up approach to distributed or cooperative problem-solving in specific domains, based more or less implicitly on human models², the research directions have diverged from unifying concepts toward complex, idiosyncratic solutions.

The goal of DAI research has been the development of general problem-solving heuristics for cooperative interactions among computers. It is one of two major lines of research concerned with distributed, multi-computer solutions to complex problems. Distributed processing (DP) considers the use of multiple processing components for solving problems which are readily decomposed into fixed subproblem sets [1]. DAI is concerned more with the use of logically distributed components³ to solve problems which are perceived to be too large for solution by a single component but which are not readily decomposable in a fixed form [2]. Problems which are particularly appropriate for DP are those which decompose naturally into weakly coupled or independent processes for implementation as parallel computations. DP research has emphasized the physical distribution of computing resources from both the hardware and software points of view. Although the computing requirements of DAI systems often necessitate distributed processing for implementation,⁴ research has been focused on such issues as relationships among problem-solving nodes (organizational structure and communication/connectivity networks) and the planning by a single agent of multiple agent activities. This review is concerned exclusively with DAI; related distributed processing issues will not be discussed.

1. Distributed Problem Solving (DPS) is an acronym which is often used in the literature - we use DAI as its synonym.

2. See, for example [9] [28].

3. Components of distributed problem-solving systems are variously termed agents, nodes, or knowledge sources depending on the viewpoint of the author; agent or node will be used here with no attempt to make a distinction between the two.

4. For discussions of the process of physically distributing a logically distributed problem-solving system, see [3,10]

In lieu of a traditional review, we describe three DAI systems which have been implemented in simulation and which are representative of the contexts within which DAI is currently being explored. These simulations will form a basis for subsequent discussion of the major issues addressed in the DAI literature: 1) the Distributed Vehicle Monitoring Testbed (DVMT) of Lesser, Corkill, and co-workers [3,4,5,6,7]; 2) the Distributed Sensing System (DSS) of Smith and Davis [8,9]; and 3) the RAND Remotely Piloted Vehicle (RPV) simulations [10,11]. The first two simulations model the interpretation of signals by networks of stationary sensing elements, while the third models missions of autonomous aircraft in benign and hostile environments.

The DVMT is a simulated distributed sensing network designed for empirical studies of control organization and communication policies. This system was based on a model of distributed problem-solving derived from the Hearsay-II system for natural language understanding [12], in which a solitary problem-solver is comprised of a number of specialized knowledge sources linked through a centralized database, the blackboard. In the DVMT, sensors are distributed over an area containing genuine and spurious vehicle tracks; although modeled after a system of specialists, the nodes (sensors) typically share the same basic skills. Tracks are stationary but extend beyond the sensing range of any single node. The network task is to arrive at descriptions of genuine tracks, requiring the exchange and integration of partial results from different sensors. Systems of this type are often referred to as "result-sharing", to distinguish them from "task-sharing" systems [8] (see below).

The Distributed Sensing System of Davis and Smith consists of a simulated network of stationary nodes, each of which can sense, process sensory data or integrate results. It is a network of specialists. The network task is to identify vehicles from simulated acoustic signals and track them during vehicle motion. The DSS is designed to explore the utility and implications of a specific protocol for organizational structuring and communication, the Contract Net [8]. The Contract Net was designed with groups of human experts working together as the model of cooperation and is intended for problems which decompose naturally into hierarchically organized subproblems. The Contract Net protocol specifies a form of negotiation to be used in matching agents with sub-tasks and assembling subproblem solutions into global solutions. It therefore offers a mechanism for distributing subproblems and their solutions among problem-solving agents; for this reason, the DSS is termed a task-sharing system.

RAND's ongoing program of developing simulations of RPV's (Remotely Piloted Vehicles) grew out of their simulations of distributed Air Traffic Control [11,13,14,15], in which individual agents were each responsible for controlling a designated aircraft's flight in a congested air space. For the RPV simulations, RAND added simple surveillance missions to be executed by groups of autonomous aircraft in different environments. Both the DVMT and DSS are intended to address problems of distributed interpretation⁵ rather than distributed action. In each case, interaction among agents occurs strictly through communication. Although the DSS involves the interpretation of sensor input from a dynamic environment, nodes are not capable of altering the environment in any way. The RAND work differs from the DVMT and DSS simulations in that agents are allowed not only to observe the environment but to move within and alter it; as a result, agents interact not only via communication but also through their effects on the problem environment.⁶ In this respect, RAND has moved closer to investigating the requirements of cooperative agents capable of action in the real world.

Coherence.

A theme which can be traced through the work described thus far, and through much of the DAI literature, is coherence, i.e. how the activities of separate agents may be integrated into a single problem solution.^{7,8} The notion of coherence in DAI closely parallels that of means-end coherence [16], the requirement that elements of a single agent plan be reasonably expected to lead to the desired end. However, the coherence problem is treated in the DAI literature as if it results from the fact of multiple agents rather than single agents. It has typically been cast as the problem of obtaining globally relevant action from agents which have only local knowledge [2,4].⁹ In spite of this point of view, some authors have recognized that

5. Distributed interpretation can be viewed simply as the integration of information such as sensor data from distributed sources or interpretation by a distributed system of information which may or not come from a single source. The term data fusion is commonly used for the former, whether or not the interpretation is carried out in a distributed system [c.f. 10].

6. Interactions among agents through their effects on the problem were also noted in the DVMT studies in the form of agent distraction by the exchange of partial problem solutions [6].

7. "Coherent cooperation is the holy grail of distributed problem solving network research." [7, p. 43.]

8. "In summary, a main challenge to distributed problem solving is that the solutions which ... must not only be reasonable with respect to the local task, they must be globally coherent and this global coherence must be achieved by local computation alone." [14, p. 767]

9. Davis and Smith cite "the fundamental conflict between the complete knowledge needed to ensure coherence and the incomplete knowledge inherent in any distribution of problem solving effort" [9, p. 64].

the use of multiple agents not only adds problems but also solutions¹⁰, but little effort has been devoted to analysis of the relationship between the coherence issues in single agent and multi-agent problem-solving. Diverse problems have been subsumed under the heading of coherence including interference among agents [17] and the efficient use of problem-solving resources [4].

Based upon a paradigm of human experts as problem-solvers,¹¹ initial attempts to ensure coherence have relied on centralization of control, supported by extensive communication. Attempts have been made to study the role of organizational structure in distributed problem-solving from a general point of view [c.f. 4,18], the only models considered having been those used for human socio-economic organization¹². No theory has yet been developed to permit organizational structures to be matched with problem characteristics. Existing DAI systems have adopted specific ad hoc structures which originate from human models in particular domains [9,10,19]. Of necessity, communication and control have been considered simultaneously in the analysis or design of organizations,¹³ although the two are frequently separated for the purposes of discussion. As a consequence of the use of human models, the assumption has been made that hierarchical structures will, in general, require less communication than less centralized control structures.

The RAND RPV simulations started with a fixed hierarchical organization [11].¹⁴ Empirically, RAND found that optimal group performance could be obtained with such a structure only in static environments or situations of low stress. Stress refers to such factors as uncertainty in communication, an unpredictable or hostile environment, or unreliability and failure of individual agents. A negotiation procedure was added as a means to introduce flexibility by reassigning roles after small changes in low stress environments [10]. The authors recommended that role reassignment in high stress environments be carried out according to a fixed default succession or by assignments from the group leader, as long as the

10. "...the key to coherent distributed problem solving lies in the fact that while distributed agents have greater difficulties in solving a given task, they have potentially more options as well.... In short, much of the power of distributed problem solving comes through cooperation and communication" [14, p. 767].

11. Examples of such paradigms include the cooperating experts metaphor of [9] and the Scientific Community Metaphor of [29].

12. Consider, for example, the stated assumption in [18] that "modules, processes, and tasks act like humans in an organization" (p. 79).

13. For example, pre-defined lines of communication can establish de facto lines of control and control structures must be supported by compatible communication policies.

14. For a discussion of fixed and flexible hierarchies see [30, Section 4].

status of individual agents (RPV's) was only minimally changed. However, they offered no recommendations for mechanisms by which to respond to large scale changes in RPV status during high stress situations. In general, the effectiveness of a fixed hierarchy dropped with increasing stress and the authors suggested that more "robust distributed organizations [will] become favored", although the forms such organizations might take were not specified.

The DSS of Davis and Smith is organized as a flexible hierarchy, in which the roles played by individual agents may vary during problem-solving, but the overall hierarchy is maintained. Davis and Smith have pointed out that negotiation is an expensive protocol and that increased numbers of communication connections are required in a flexibly organized network as compared to a fixed hierarchical structure. Their conclusion is that the communication and computation overhead involved in a flexible hierarchy is justified only when subproblems are large. For appropriate problems, i.e. those with natural hierarchies of task or of data abstraction, they assert that the potential for increased reliability and resource utilization in flexible rather than fixed hierarchies warrants the expense of negotiation. It should be noted, however, that Davis and Smith considered flexible hierarchies only in relation to fixed hierarchies, and not to less centralized structures.

A distributed problem-solving system built by Yang, *et al.* [20], for designing digital-logic circuits is structured as a flexible hierarchy, similar to the DSS. The design task is initially given to a single node, which carries out problem decomposition, makes subproblems available to other nodes and is responsible for synthesizing results. The organization of this system is similar to the Contract Net, except that the originating node has no control over which individual node(s) work on subproblems except by directing subproblem descriptions to individuals. No negotiation protocol is involved.

A formally decentralized control structure was employed in most of the DVMT experiments, although the effect of introducing a specialist with responsibility for integrating partial results was investigated briefly [6]. Track signals were placed in different positions with respect to the sensors, creating problems which decomposed along different lines. Implicit organizational structure was varied by externally controlling the relative weight given by individual nodes to self-generated and externally-generated data and results. When results from a number of scenarios were compared, no single organizational structure was found which

allowed the network to function optimally under all conditions.¹⁵ Optimality in one scenario was often associated with poor performance in another. Noting that network reorganization can be costly and time-consuming, the authors concluded that organizational structures will have to be chosen carefully for long-term performance in real systems.¹⁶

Communication as a mechanism for enhancing coherence by increasing the knowledge of individual agents, both in the problem domain and about each other has been an essential component of all of these simulations. An important practical question that has arisen in the DAI literature is how to minimize communication loads without sacrificing functionality. Two approaches to communication reduction have been taken: limiting message circulation (the number of recipients) and limiting message content (the type of information to be communicated), thereby reducing the number or length of messages. Both approaches require consideration of the forms of communication (language, syntax) in specific problem domains,¹⁷ although this will not be discussed further in this review.

The two extremes of a circulation mode are represented by broadcast (communication to all potential receivers) and directed transmission (communication to a single receiver). Of the two, broadcast is the most commonly implemented form in the DAI literature. The DSS and DVMT incorporate only broadcast. Huhns, et al., implemented a hierarchical communication network with broadcast occurring within levels and directed transmission between levels. Yang, et al. [20], in the logic design testbed, implemented both broadcast and directed transmission and enabled the transmitter to select between the two, based on the extent of available knowledge about potential receivers. In the RAND RPV simulations, communication between agents required multiple steps because the distance between certain aircraft could exceed the communication range. Therefore, two modes of directed transmission were implemented: transmission along a fully preplanned route of intermediate agents ("routing table") and by a sequence of individually planned steps ("chain letter"). In addition, true broadcast and non-selective re-transmission was employed ("spreading activation"). The authors found that the utility of each circulation mode depended upon the degree of situational stress; proper selection required accurate assessment of conditions.

15. Performance was judged by the number of communication acts, total time and spurious hypotheses generated en route to the solution.

16. The DVMT dealt with a static problem; dynamic problems will require systems which can operate effectively, if not optimally, under a variety of conditions.

17. c.f. [20, 22].

Durfee, et al. [7], in their work with the DVMT concentrated on communication content rather than on circulation. They examined the effectiveness of communicating information at different stages of the problem-solving process: early (incomplete) results, locally complete results¹⁸ or a mixture of the two. Their work revealed a complex interaction between the type of information transmitted, the weight given to received information by the receiver, and specific problem characteristics in determining system performance. Although "judicious" communication of partial results was shown to decrease overall communication requirements, no single communication policy proved to enhance system performance in all circumstances.

Frequent reference is made in the DAI literature to an intimate connection between control structure and communication.¹⁹ The use of hierarchical organizations has been defended as a means to reduce communication requirements by limiting circulation. In practice, however, the work discussed above [9,10] has demonstrated that extensive communication is required to support hierarchical organizations. At this time, the tradeoffs between centralization and communication load remain incompletely understood. An overall theory is lacking by which to determine the relative control requirements and corresponding communication needs of distributed solutions to different types of problems. Research has proceeded by selecting particular (primarily centralized) organizational structures, adding flexibility in response to empirically demonstrated needs, and tailoring communication policies to minimize overhead in the resulting systems.

Planning: a Single or Multiple Agent Task?

In accord with the emphasis on centralized organizational structures, cooperative planning in the DAI literature has primarily referred to centralized planning, that is, planning by one agent for the actions of multiple agents.²⁰ The hope has been that centralized multi-agent planning systems could be developed using as a basis previous work on single agent planners.²¹ The main thrust of planning research in DAI has therefore been to identify the unique aspects

18. Locally complete results are those which are estimated by a node to be the best it can obtain itself.

19. In fact, Fox [18, p. 71] groups control regime and communication paths into "organisational structure".

20. The planning requirements of fixed and flexible hierarchies are similar, with the exception of decisions regarding which agent is to have control, although in the first case centralization is permanent and in the second it is temporary.

21. Nilsson has turned this around by suggesting that solutions to multi-agent planning problems may advance the field of single agent planning: "Methods used by one AI system for reasoning about the actions of other AI systems will also be useful for reasoning about other dynamic (but unintelligent) processes in the environment." [31, p. 43]

introduced by having a single agent reason about the actions of multiple agents rather than reasoning only about its own. Elements which have been associated with centralized planning include the ability to represent concurrent events, to identify spatial and temporal relationships among concurrent events, especially potential interference between the actions of multiple agents, and to represent and reason about the actions of other agents. Although certain of these elements, particularly the last will also be important in decentralized planning, centralized planning implies reasoning about the actions of agents *over which the planner has control*. Therefore, solutions to the centralized multi-agent planning problem will not necessarily provide solutions for distributed planning, in which other agents may be more or less predictable but whose actions cannot be established in advance.

Georgeff [17], Lansky [21] and Stuart [22] have each proposed mechanisms for developing and representing multi-agent plans. They have defined methods to represent explicit temporal constraints among actions and rules by which to apply these constraints in the construction of a plan. In each case, concurrent plans are constructed by reducing actions to primitive components and interleaving the components of different tasks in a single sequence. The focus of these efforts is on the relationships among actions or behaviors, rather than on sequences of world states, which led Georgeff and Lansky to propose a behavior-specification or procedure-based approach to planning in general [23]. The purpose for developing formalisms for handling these particular aspects of centralized planning was so that they could be integrated into existing single agent planning frameworks to form multi-agent planners. However, we are not aware of any distributed systems which have yet incorporated the results of efforts along this line.

Konolige and Nilsson [24] have been concerned with the way agents represent and use beliefs about each others capabilities and actions. Konolige and Nilsson's approach, suitable for either centralized or decentralized planning, concentrates on ways by which one agent could take into account the probable actions of others while developing its own plan.²² Their solution was based on an assumption that all agents' actions result from generating and executing plans to achieve specific goals. When their goals and planning processes are known, the actions of other agents can be predicted and incorporated into a plan. Questions of spatial and temporal coordination were not considered by Konolige and Nilsson. Rosenschein [25] also proposed ways to represent beliefs about other agents specifically for the purpose of a centralized planner. In addition, Rosenschein also suggested a syntax for expressing constraints in

22. Konolige and Nilsson did not assume that all goals were the same but only that they were known.

commands to subordinate agents which distinguished between those which were necessary and those which were only desirable. By permitting the central planner to alter the belief and goal states of other agents, Rosenschein also specified a mechanism for exerting control. Durfee, et al., initially incorporated only solitary planners in the DVMT. A form of distributed planning was introduced at a later time, allowing nodes to access each others' top level plans and goals. No centralization was introduced by this approach but the ability of nodes to coordinate their separate problem-solving activities was enhanced.

Several forms of planning to determine individual and group actions were considered in the RAND RPV simulations. All planning was carried out by a single agent, the "leader", for the entire group. As the RPV simulations became more complex, RAND identified three necessary types of planning: rule-based planning, simulation-based planning and logic programming. In situations of high stress and limited knowledge (i.e. increased uncertainty), condition-action rules were used to select one action at a time; this rule-based planning resulted in a series of individually selected discrete actions. Simulation-based planning was found to be useful in low-stress environments with few alternative sequences; a single sequence of actions was selected from among a suite of alternatives by simulating their execution in advance and comparing the anticipated consequences.²³ To select from among a number of options while satisfying specific constraints, the authors recommended logic programming, building a multi-step plan from a set of primitive actions.²⁴

Agent Models and Cooperative Heuristics.

There is common agreement among researchers in the field that robust and coherent distributed problem-solving will require more flexibility than is possible within current approaches. In a single static problem domain, the DVMT, no single organizational structure or communication policy has been found to be equally effective²⁵ for different problem variants. The RAND RPV simulations demonstrated that flexibility of organization and communication (with correspondingly less centralization) is a necessity in dynamic

23. This is clearly a costly activity; RAND acknowledges that this form of planning is only appropriate in situations that "require deep planning" without indicating how such situations should be identified.

24. RAND noted that logic programming requires "substantial situation knowledge", but did not specify further conditions. Presumably, logic programming requires more computation than rule-based planning and would therefore be appropriate only in low-stress environments.

25. See note 15.

environments. In addition, the computational and communication overhead associated with fixed and flexible hierarchical organizations has become increasingly difficult to support in complex systems and environments. As a way to introduce flexibility, reduce overhead and increase coherence, DAI researchers have expanded the kinds of knowledge shared among agents. A basic level of knowledge about distributed problem-solving implicitly shared among agents was incorporated in DAI research at the outset, the benevolence assumption.²⁶ This is the assumption that agent actions are directed to the same goal or at least that agents will refrain from interfering with each other. With time, the need for additional knowledge, either built-in or acquired during problem-solving, has been demonstrated empirically in each of the DAI systems described above. The emphasis in the literature has been primarily on the sharing of information about agent capabilities.

In their distributed system for designing logic circuitry, Yang, et al. [19,20], provided a mechanism to associate subproblem descriptions with lists of agents capable of offering solutions. These lists were updated during the problem-solving process, a simple learning scheme, and were used to limit the number of agents to whom subproblem descriptions were circulated. In the DSS, agents make explicit statements during the negotiation process concerning their own potential to solve subproblems. These statements were then used according to fixed rules to determine subproblem distribution. In decision-theoretic approaches like that of Ginsberg, et al. [26,27] as well as in the work of Konolige and Nilsson [24], knowledge about agent goals and payoffs is a prerequisite for the selection of individual actions or development of plans. Decision rules or rationality constraints then determine successive actions to be selected. In the DVMT, nodes were given direct access to certain of the results and problem-solving plans of other agents for use in forming their own plans. Two kinds of knowledge about agent capabilities were incorporated into the RAND RPV simulations, one suitable for projecting future courses of action and the other for inferential reasoning [10]. This followed closely the earlier suggestion of Konolige and Nilsson [24] that knowledge about other agents be represented in two ways: by modules which reproduce the reasoning processes of those agents and by explicit beliefs about other agents. In the RAND simulations, simple rules were used for role reassignment, although other mechanisms have been planned [10, p. 38].

These approaches have focused on the sharing of information about the problem domain, of what agents *know*, and knowledge about agent capabilities (what other agents can or will *do*). We refer to the latter as agent models. While this term implies something more extensive than

26. Notable exceptions are [27, 24].

has typically been discussed and implemented, it describes the concept for which researchers appear to be reaching, *viz.* the ability to predict and account for the actions of other agents during problem-solving.²⁷ To date, agents have been provided with only a few simple rules for using knowledge about each others' capabilities and plans. It seems that no systematic study of agent models and heuristics for their use in problem-solving has yet been undertaken. In fact, the distinction between these two forms of knowledge has not even been entirely clear in the literature.

Specificity at the Cost of Generality.

The variety of implementation approaches and wealth of detail in the literature obscure the central goal of DAI research--to develop concepts and general heuristics for distributive or cooperative problem-solving. Is this goal being achieved; have concepts been developed which lead to a more generalized understanding of distributed problem-solving? It appears that the answer must be a qualified no. Although discrete issues are being addressed, such as the value of negotiation protocols and mechanisms for implementing them, no organizing principles have emerged from the work done to date. Given the resources and talent devoted to this research, why has the goal remained elusive? The answer to this question may lie in way the research has been carried out.

Many examples of the traditional approach that has been taken in DAI may be cited. The specific goal of one important example research program at RAND has been to develop a set of "operational guidelines" for executing certain RPV missions. They have analyzed specific RPV (or ATC) problems (missions), and examined architectures and protocols for solving them.²⁸ This work has been more extensive than most of the published DAI research. However, initiating research by consideration of specific problems or applications entails the risk that results will be of limited applicability; without generality, solutions to related problems may have to be derived by starting anew. This bottom-up approach, beginning with particular problems and applying specific (human) models, has been characteristic of DAI research. As the effectiveness of the models and their initial assumptions has broken down, considerable

27. The expressions "metaknowledge" and "metalevel knowledge" have frequently been used in the literature to refer explicitly to shared knowledge about agent capabilities [cf. 7, 20]. We have avoided these terms in this context to prevent confusion with other uses.

28. This work is ongoing and the researchers recognise that it is early to generate guidelines even for specific applications. [10, p. 41].



effort has been expended on incremental, compensatory improvements. This has led to complexity and confusion and suggests the value of a higher level view of distributed problem-solving in order to make significant advances in the field and to consolidate gains made thus far.

REFERENCES

1. R. G. Smith, Report on the 1984 Distributed Artificial Intelligence Workshop, AI Magazine, 234(Fall, 1985).
2. Randy Davis, Report on the Workshop on Distributed AI, IEEE SIGART Newsletter 73, 42(1983).
3. Edmund H. Durfee, Daniel D. Corkill and V. R. Lesser, Distributing a distributed problem solving network simulator, Proc. Real-Time Systems Symposium 1984, 237(1984).
4. Victor R. Lesser and Daniel D. Corkill, Functionally Accurate, Cooperative Distributed Systems, IEEE Trans. on Sys., Man, and Cybern. SMC-11, 81(1981).
5. V. Lesser, D.Corkill, J. Pavlin, L. Lefkowitz, E. Hudlicka, R. Brooks, and S. Reed, A high-level simulation testbed for cooperative distributed problem solving, IEEE 3rd Int. Conf. on Distributed Computing Systems, 341(1982).
6. Daniel D. Corkill and Victor R. Lesser, The use of meta-level control for coordination in a distributed problem solving network, IJCAI 83, 748(1983).
7. E. H. Durfee, V. R. Lesser and D. D. Corkill, Coherent cooperation among communicating problem solvers, Draft, (September, 1985).
8. R. G. Smith and R. Davis, Frameworks for cooperation in distributed problem-solving, IEEE Tran. on Sys., Man and Cybern. SMC-11, 61(1981).
9. Randall Davis and Reid G. Smith, Negotiation as a Metaphor for Distributed Problem Solving, Artificial Intelligence 20, 63(1983).
10. R. Steeb, D. McArthur, S. Cammarata, S. Narain, J. Rothenberg and W. Giarla, Cooperative intelligence for RPV fleet control: analysis and simulation, RAND Note, Working Draft WD-2598-ARPA, (April, 1985).
11. R. Steeb, D. McArthur, S. Cammarata, S. Narain, and W. Giarla, Distributed problem solving for air fleet control: framework and implementation, The Rand Corporation N-2139-ARPA, (April, 1984).
12. L. D. Erman and V. R. Lesser, A multi-level organization for problem solving using many, diverse, cooperating sources of knowledge, Advance Papers of the 4th IJCAI, 483(1975).
13. D. McArthur, R. Steeb and S. Cammarata, A framework for distributed problem solving, Proc. Nat. Conf. Art. Intell., 181(1982).

14. S.Cammarata, D. McArthur and R. Steeb, Strategies of cooperation in distributed problem solving, The Rand Corporation N-2031-ARPA, (1983).
15. R. Steeb, S. Cammarata, and S. Narain, Distributed intelligence for RPV fleet control, Proc. IEEE EASCON , 223(1984).
16. Michael Bratman, Taking plans seriously, Social Theory and Practise 9, 271(1983).
17. M. P. Georgeff, A theory of action for multiagent planning, Proc. 4th Nat. Conf. on Artificial Intelligence , 121(1984).
18. Mark S. Fox, An organizational view of distributed systems, IEEE Trans. Sys., Man and Cybern. SMC-11, 70(1981).
19. Michael N. Huhns, Larry M. Stephens, and Ronald D. Bonnell, Control and cooperation in distributed expert systems, IEEE Southeascon , 241(1983).
20. J-Y. D. Yang, M. N. Huhns and L. M. Stephens, An architecture for control and communications in distributed artificial intelligence systems, IEEE Trans. Sys., Man and Cybern. SMC-15, 316(1985).
21. A. L. Lansky, Behavioral Specification and Planning for Multiagent Domains, SRI Tech. Note 360 , (1985).
22. Christopher Stuart, An implementation of a multi-agent plan synchronizer, Proc. 9th Int. Joint Conf. on Art. Int. , 1031(1985).]
23. M. P. Georgeff and A. L. Lansky, Procedural Knowledge, Proc. IEEE 74, 1383(1986).
24. Kurt Konolige and Nils J. Nilsson, Multiple-agent planning systems, Artificial Intelligence Conference Report SRI International, 138(1980).
25. Jeffrey S. Rosenschein, Synchronization of multi-agent plans, Proc. Nat. Conf. Art. Int. , 115(1982).
26. M. R. Genesereth, M. L. Ginsberg, and J. S. Rosenschein, Cooperation without Communication, AAAI Conf. Proc. 1, 51(1986).
27. Matthew I. Ginsberg, Decision Procedures, DTIC Report STAN-CS-85-1064, 1(1986).
28. S. A. Cerri, P. Landini and J. Leoncini, Cooperative agents for knowledge-based information systems: dialogue about the archeology of Rome, App. Artificial Intelligence 1, 1(1987).
29. C. Hewitt and B. Kornfeld, Message passing semantics, IEEE SIGART Newsletter 73, 48(1983).
30. R. F. Franklin and L. A. Harmon, Elements of Cooperative Behavior, IR&D Final Report, Draft , (February, 1987).
31. N. J. Nilsson, Two heads are better than one, IEEE SIGART Newsletter 73, 43(1983).



APPENDIX B

The Nature of Cooperation in Humans and Animals

A Report for

The Environmental Research Institute of Michigan

Robert J. Taylor
Department of Fisheries and Wildlife
Utah State University
Logan, Utah 84322-5200

INTRODUCTION

This is a review of the basic literature on the subject of cooperation in humans and animals. Depending upon how one chooses to define cooperation, this literature ranges from minuscule to massive. If one considers only highly coordinated behaviors of general importance, the published research is quite small. If, on the other hand, one considers cooperation to be any behavior in a social context, then the literature is staggering. Anticipating that the primary use of this review will be to help discover metaphors for machine cooperation, I have chosen to concentrate upon that fraction of the literature which yields general lessons. For example, the basic work on cooperation in humans is published in the social psychology literature; however, applications appear in fields as wide-ranging as business, education, law, sports, and the military. The focus in these applied areas is almost always quite narrow; little, if any, attention is directed to discovering the basic elements of cooperation.

Throughout the main body of the review, I avoid reference to cooperation and coordination in machines. The focus is entirely upon what is known about cooperation in man and animals. In the conclusion I shall hazard a few opinions about what all this means for machine intelligence.

The explicit study of cooperation is relatively new. Students of both human and animal behavior seem to have discovered the topic about 20 years ago, and the majority of all work has been done in the last ten years. This notwithstanding, it is a popular area of research, and the pace of publication is accelerating.

The best place to begin a review is with a definition of the term of interest. As is true of important terms in general, people use different definitions of cooperation; unlike with other terms, scientists seem remarkably unconcerned about this. Social psychologists have, without exception, adopted a narrow, functional definition. To them, cooperation is observed only in highly structured experimental situations; as a consequence, they define the behavior entirely in terms of the reinforcement contingency (Schmitt 1984). If a group is rewarded as a whole for its performance on a task, then the behavior that results is cooperation. If individuals receive separate rewards, then this circumstance will result in individualistic or competitive behaviors. Note that the actual behaviors are not part of the definition. The reason for this is that the

suite of behaviors that experimental subjects can employ is invariably limited, usually quite severely, by the design of the experiment. The question of what comprises cooperation is determined *a priori* entirely by the experimenter.

Animal behaviorists rely much more heavily upon unstructured observations of animals in their natural environments. One would expect, therefore, that they might have devoted more attention to discovering precisely what cooperation consists of. This turns out to be true, to some extent. But these people face another problem; the enormous variety of circumstances in which cooperation arises and the special features of each circumstance make any single definition quite local to the animals and environment of that study. The consequence is that while each animal behaviorist has a fairly good idea of what she means by cooperation, people seem unprepared to assert a general definition. I shall return to this question of definition later in the review.

A classification of published research on cooperation can employ, with equal validity, either of two alternative dichotomies. The first, and most obvious, is to divide the literature into studies of animals and studies of humans, and I shall follow that division here. It is important to note that this difference goes beyond the organism studies; it reflects at a fundamental level the sorts of questions scientists have asked. Social psychologists strongly emphasize rational and social explanations of cooperation. This accounts both for their heavy reliance upon the theory of games for models and for their interest in repeating the same kinds of experiments in an apparently endless variety of social contexts. The consequence is that human behaviorists strongly emphasize, perhaps overemphasize, the plasticity of this class of behaviors. Animal behaviorists, by comparison, are strongly biased to favor genetic, evolutionary mechanisms and tend to assume that the nature and degree of cooperation in a single animal is fixed by its genetic heritage. Their research is heavily comparative in nature.

A less obvious dichotomy, although equally valid, is between scientists who ask whether cooperation should occur in a particular circumstance and those that ask how, given that it occurs, it should proceed. These separate questions arise both in animal studies and human studies, although they are given somewhat different emphases in the two. For example, psychologists never seem to question whether cooperation occurs at all in humans; its existence is assumed to be obvious. The question of cooperation's actual existence is of primary importance in animal behavior and must be answered at the outset of any research project. Nor is the question of why humans should cooperate often asked by psychologists; some tasks seem

to lend themselves, without dispute, to cooperative solutions. Again, this is less obviously true of animals.

These differences aside, all fields of research share a profound interest in the question of whether a subject will or will not cooperate in any particular set of circumstances. This is, in fact, the overwhelmingly dominant theme in research on cooperation. Game theory is the major conceptual approach to this question in both animal and human studies. The theory is usually posed in such a way as to grant clear, discrete alternative strategies to the human or animal subject: cooperate with others, ignore them, or compete with them.

How one is to cooperate is an entirely different question from whether to cooperate and has not attracted nearly as much attention. Perhaps this should not be surprising in the area of human behavioral research, given the predilection of social psychologists for conducting highly simplified experiments. Cooperation versus competition may be equivalent to nothing more complicated than pushing button A versus button B. The potential richness of cooperative behaviors that could emerge in less-structured circumstances seems to be of relatively little interest to students of human behavior.

One would expect the opposite to be true of animal behaviorists, that their observations of animals in natural surroundings would emphasize more strongly the mechanics of cooperation. This does not appear to be the case, however. The reasons, I suggest, are two. The first is quite obviously theory intoxication. Game theoretical models entered into the social sciences quite early, in the 1950's and 1960's, but animal behaviorists did not discover them until the mid-1970's. The conceptual richness of these models has subsequently overwhelmed other aspects of research on cooperation. To some extent this is a fad, but in defense of biologists, the selfish rationality of the game-theory perspective appeals to the modern worker. This brings me to the second reason underlying the emphasis on why rather than how. After the intellectual poverty of the group-selection explanations of altruism and cooperation was revealed in the mid-1960's, biologists embarked on a crusade to purge all such fuzzy thinking from their midst. The dominant attitude at present is perhaps best expressed by Ghiselin (1974):

No hint of genuine charity ameliorates our vision of society, once sentimentalism has been laid aside. What passes for cooperation turns out to be a mixture of opportunism and exploitation. The impulses that lead one animal to sacrifice himself for another turn out to have their ultimate rationale in gaining advantage of a third; and acts "for the good" of one society turn out to be performed to the detriment of the rest. Where it is in his own interest, every organism may reasonably be expected to aid his fellows. Where he has no alternative, he submits to the yoke of

communal servitude. Yet given a full chance to act in his own interest, nothing but expediency will refrain him from brutalizing, from maiming, from murdering - his brother, his mate, his parent, or his child. Scratch an "altruist," and watch a "hypocrite" bleed. (p. 247)

Given this perspective on animals, one can more easily appreciate how the discovery of selfish explanations of acts of apparently cooperative altruism can provide an entertaining avenue of research. In spite of this, one would expect some biologists to have chosen the less-traveled way, to be looking at the actual mechanisms of cooperation. That so few have done so can only be attributed to fashion.

COOPERATION IN HUMAN BEHAVIOR

The nature and extent of cooperation among humans is a subject of active research at present. Why this should be the case is not difficult to discover. At the most fundamental level, the question of whether we are by nature cooperative cuts close to the heart of our conception of ourselves (Schumacher 1984). Our political system must inevitably mirror those conceptions, those beliefs about human nature. Marx, for example, observed that the whole idea of liberty presupposes that we find in other people the limits of our freedom, not the realization of it, an assumption with which he, and many others, took issue.

On a more mundane but still important level, it is of considerable practical value to us to discover how humans cooperate. We are educated in groups, we work in groups, and we fight in groups. That these group activities be done well is clearly important to a healthy and secure society.

Perhaps the best place to begin examining cooperation in humans is to review social variation in the phenomenon. McClintock and Van Avermaet (1982) have erected a taxonomy of social values, using a coordinate system in which the abscissa is one's own benefit and the ordinate is the benefit of others (see Figure 1). Only those acts leading to at least some positive benefit for the actor are considered to provide a stable basis for the formation of a society.

This leads to a classification of societies, taken from the right half of the coordinate system, in which social systems are arranged along an axis from highly cooperative to highly competitive, with individualistic societies at the midpoint. This attempt at definition appears to be acceptable to most social scientists (although biologists define altruism quite differently, *sensu* Hamilton (1963)).

Bethlehem (1982) reviewed the extensive anthropological literature on cooperation to discover examples of these three types among non-Westernized societies. Perhaps the best examples of cooperative societies are the Bushmen and Mbuti (pygmies) of Sub-Saharan Africa. These peoples are unconcerned with social position or wealth and identify themselves almost entirely with the group. Asian societies are characteristically competitive. The Ifugao of the Philippines, for example, seem to be engaged in a constant struggle for relative social position.

Modern Western societies fit, on the whole, into this category. The classic example of an individualistic society is the Ik culture of East Africa, whose members are amazingly oblivious to one another. They share nothing, even within families. They engage in no group activities. Children beyond the age of weaning are not fed; they survive by literally stealing food from the mouths of the aged. The Ik are considered anthropological curiosities, the result of short-term adaptations to a very poor environment and with little potential for survival as a social unit. But the more extreme cooperative and competitive societies are not. These are alternative forms of a stable social order, and their existence alone gives pause to those who would proclaim human nature either fundamentally cooperative or fundamentally competitive.

Cross-cultural studies of children from modern societies suggest that those raised in environments which emphasize cooperation act, by-and-large, more cooperatively. Rural and traditional societies and Israeli kibbutzim, for example, give rise to cooperative children; whereas, urban, Westernized societies do not. There is a tendency to cooperate more with individuals from within the parent group than with those from without (Espinoza and Garza 1985).

Very young children are usually unwilling to cooperate at all, except perhaps in simple ways with their parents (Lawrence 1984). This is related to linguistic maturity rather than age. The appearance of cooperation in humans, therefore, requires a minimum amount of cognitive and communicative development. Cooperation increases in frequency until ages 8-11 in American children and may actually diminish somewhat beyond that point, particularly in boys (Stingle & Cook 1985). Sex differences are not apparent among small children (Georgas 1985a, 1985b) but may emerge in older children (Stingle and Cook 1985, Stockdale, et al 1983). Cooperation remains high among girls, and girls' performances may suffer relative to boys' when they are forced into competitive situations (Johnson, et al. 1985).

There is ample evidence that the perceived social situation dramatically alters the tendency to cooperate. To some extent people have an "interpersonal style" that determines how they will interact (McLeod and Nowicki 1985), but this can be modified by experience. Children may become more cooperative in games when faced with a cooperative partner (Brady, et al. 1983). Adults playing the prisoner's dilemma game appear quite willing to change their behavior back and forth between cooperating and defecting (cheating), following the lead of their opponent (Lindskold, et al. 1986).

In these games, communication is important for maintenance of cooperative responses (Braver and Wilson 1986, Liebrand 1984, Stults and Mees 1984). When players cannot communicate, they assume the worst of their partner and switch from cooperation to competition. The same things occurs in rats (Daniel 1942, 1947; Gardner, et al. 1984).

Within-group cooperation appears quite sensitive to the presence of other competing groups (Rabbie 1982). A well-known study by Blake and Mouton (1961) of business managers put into conflict situations discovered that groups in intense competition with other groups go through a predictable sequence of changes. Group members close ranks; they become more cohesive and consider themselves superior. The group becomes more hierarchical; centralized leadership becomes acceptable. Loyalty and conformity are demanded of the members, and deviating opinions are not tolerated.

One must accept that these studies reveal just the tip of the iceberg, that a variety of other factors influencing cooperation remain to be discovered. For example, Abric (1982) reported the results of college students playing the prisoner's dilemma game against a computer. Half the students were told they were playing against a computer, and half were told they were playing against another student. The computer's responses were the same in both cases. Those who thought they were playing against the computer cooperated very little. Those who thought they were playing another student cooperated a great deal. Halfway through the game the students were informed that the experimenter had lied, that the opponent was really the opposite. Their tendency to cooperate immediately reversed itself. This does not appear to be terribly rational, but the students apparently perceived a human opponent as potentially more cooperative. The machine was assumed to be rigid, objective, and unresponsive.

Anthropologists excepted, social scientists have traditionally inferred these patterns from the results of experimental games played by subjects in controlled surroundings. The history of this methodology is reviewed by Colman (1982). The origin of experimental game research lies in the application of the theory of games to human behavior. Solutions to theoretical games derive from the assumption that the players are dispassionate and rational and seek only to optimize their payoff from the game. An advocate of this traditional approach, Deutsch (1982), argued that

In a task-oriented relationship one is oriented to making decisions about which means are most efficient in achieving given ends. This orientation requires an abstract, analytic, quantifying, calculating comparative mode of thought in which one is able to adopt an effectively neutral, external

attitude toward different means in order to be able to make a precise appraisal of their comparative merit in achieving one's ends. (p. 32)

While the games approach has been an extraordinarily valuable heuristic theory for studies of human cooperation, criticisms of this style of research are now beginning to surface. Derlega and Grzelak (1982) identified game-theory as merely one of two separate approaches to the study of cooperation; the second is an internal, cognitive and emotional approach. Grzelak (1982) claimed, moreover, that game theory alone has not fared too well in predicting behavior in interdependence situations. His appraisal of the theory revealed three weaknesses:

1. the theory oversimplifies human motivations by assuming that people seek only to maximize their own gains;
2. people may have quite different goals and maximization rules than those assumed by the basic theory;
3. a person's social knowledge and perception of the actual degree of interdependence influence his game-playing behavior.

He made a convincing case that humans have preferences not only for larger rewards but also for various distributions of rewards and power among group members. In some cases subjects will perform "suboptimally" in order to redress what they perceive as inequities or imbalances within the group (Surazka 1986). It becomes quite clear in reading this literature that both the social context in which the game is played and the perceptions of the players are important influences upon the nature and extent of the cooperation.

Whether or not one cooperates should depend, at least in part, upon whether groups are more effective at tasks than individuals. This is obviously true in circumstances where a task is physically beyond the limits of one person, but is it always the case? Schmitt (1984) asserted that "cooperation typically leads to superior performance when task performance is facilitated by coordination, division of labor, or assistance, because these activities are reinforced only under cooperative contingencies." In learning situations groups appear to out perform individuals on average (Laughlin and Futoran 1985, Pigott and Heggie 1986). But the "on average" is important, because groups do not out-perform the best individuals; they merely evaluate individually-suggested alternatives effectively enough to recognize and adopt the best answers. One might conclude, therefore, that the least competent individuals profit the most from group learning.

Deutsch (e.g. 1982) has for years maintained that cooperative groups of people outperform the same number of individuals working alone. Modern workers no longer accept this as a general rule. I want to discuss three circumstances in particular in which group performance has the potential to exceed individual performance but often does not. The first is negotiation; the second is planning; and the third is helping.

By negotiation, I mean the circumstance where two (or more) people each have goals which they cannot achieve alone and which are to some extent incompatible. Maximizing profit on a business transaction is an example of such an interaction. This situation is one of the most complex social problems humans have to deal with (Hughes 1984). While the possibility of loss and exploitation is high, by creative and flexible bargaining the negotiators can often help one another to achieve a sufficiently large fraction of what they want that the sum of the individual benefits is well in excess of what each could achieve working independently. However, the question of whether this always occurs in group interactions must be answered no. Ben-Yoav and Pruitt (1984) identify three alternative strategies to reach agreement, yielding, contentious behavior, and problem-solving. Yielding consists of simply giving in, lowering one's aspirations. Contentious or competitive behavior consists of trying to lower the other person's aspirations. Problem-solving is trying to find a way to satisfy everyone's aspirations. Some yielding is necessary in any negotiation, but yielding farther than is necessary to achieve agreement is not good for the group. Negotiation experiments between male-female pairs of college students revealed that romantically-involved couples did rather less well than strangers, failing to identify the best strategy for the pair significantly more often. The best means of achieving maximum joint benefit was to display high initial resistance to yielding, yet be flexible and inventive about the means to achieve one's own interests, and finally to abstain from contentious behavior. Attached couples, and particularly the female members, were so strongly motivated to avoid conflict that they had a strong tendency to yield prematurely. I suspect the same would be true of subordinate members of single-sex group(s).

Problems with group planning are summarized by Janis (1982) who coined the term "groupthink." Groupthink is

... a mode of thinking that people engage in when they are deeply involved in a cohesive ingroup, when the members' strivings for unanimity override their motivation to realistically appraise alternative courses of action. (p. 9)

when group members value one another highly and those outside the group much less, they tend to make the following planning mistakes:

1. discussions are limited to an incomplete survey of the alternatives;
2. the group tends not to survey adequately the objectives to be accomplished by the planning effort;
3. in the face of new information, the group fails to re-examine its previous decisions;
4. the group neglects to consider adequately courses of action initially evaluated as unsatisfactory by the majority;
5. the group makes little or no effort to obtain expert information
6. the group exhibits a selective bias in evaluating factual information;
7. the group spends little time and effort analyzing possible setbacks and discussing contingencies.

Janis's argument is summarized in his Figure 10.1 (Appendix I). He discussed a variety of examples of groupthink in politics, ranging from the Bay of Pigs fiasco to Pearl Harbor. In each case "when groupthink tendencies become dominant, the members of an executive group withhold from each other information about their personal doubts." This phenomenon probably arises from experiences of group members with previous non-cohesive groups bogging down in unproductive clashes among adversaries holding to irreconcilable positions.

The final case of group failure is helping behavior. The Kitty Genovese case in New York, in which a woman was murdered in front of many observers, none of whom came to her assistance, sparked a great deal of research into helping behavior. This is reviewed by Latane' et al. (1981). Their conclusion is that "with very few exceptions, individuals faced with a sudden need for action exhibit a markedly reduced likelihood of response if other people actually are, or are believed to be, available to act." There is so much data available on this point that they could actually infer a general relationship between the likelihood of an individual response, I , and the size of the group, N : $I = s N^t$, where $t < 1$ and can be negative. In one study these parameters were estimated at $s = 0.77$ and $t = -.81$. Even whether the group as a whole is more likely to act than a single individual is in question. The data are ambiguous, but as much supports the hypothesis that the group is less effective than one individual than that it is more. They offer several explanations of this phenomenon, one of the most interesting of which is "diffusion of responsibility." This is the reasoning by all group members that somebody else will take care of the problem. Failure to help in these studies is probably the

consequence of poor integration of the group, a typical situation in both large cities and psychology laboratories. I have analyzed this situation using game theory in the context of alarm calling in colonial animals (Taylor, et al. in prep.) and believe the phenomenon to be evolutionarily stable.

In any of these three circumstances, negotiation, planning, and helping, groups have the potential to out perform individuals. The point of dwelling on the group failure is merely to avoid the facile assumption that groups are always superior.

The discussion of helping behaviors raises the issue of group size. If a group is superior to an individual at a task, then what size group is best? And do humans spontaneously form groups of that size? Short of working or fighting groups in modern societies, little attention has been paid to this. The usual approach is to decide what special skills are needed for a task, and let the group size follow from the sum of these. Smith (1985) has published an interesting analysis of hunting group size in Canadian Inuits. From an optimal foraging perspective, Inuits form groups that are too large for maximum profit per individual. If one incorporates genealogical relatedness and kin-directed altruism into the model, predictions of group size are improved but are still inadequate. Potential conflicts between group members and joiners and other aspects of the local social structure are apparently important as well. It is clear, at any rate, that the actual manpower needs for searching and killing comprise merely one factor in the equation determining the size of Inuit hunting groups.

If one theme emerges from this review of human cooperation, it is that the phenomenon is highly complex. People carry with them their own agendas and value systems. These can be quite subtle and frequently do not coincide with the assumptions of the researchers. A simple optimization or economic view of cooperation may be adequate in limited circumstances, particularly where subjects do not perceive themselves to have a long-term social relationship. But in circumstances where social integration is important, the game-theoretic perspective, with its attendant simplistic assumptions about motivation, seems inadequate.

COOPERATION IN ANIMALS

Individuals of most species of organisms are typically spatially aggregated in the natural setting. Whether this implies any sort of social integration or cooperation has been a subject of intense debate for decades. One would be hard pressed to justify a claim of cooperation for trees of a single species growing near one another in a stand, even though as a group they serve to modify the physical environment in such a way as to be beneficial to one another. Yet when precisely the same phenomenon occurs in animals, it is called cooperation (e.g., Chelazzi, et al. 1985). I shall restrict discussion here to cases in which coordination of behaviors toward a group goal is more explicit. This restricts animal cooperation to a small number of discrete cases: the social insects and spiders, cooperative breeding in birds and mammals, and cooperative hunting in carnivores.

Social cooperation is well-developed in many species of small arthropods. Normally it is of a fairly primitive sort, such as an increase in feeding rates in groups (e.g., Ghent 1960). Social spiders are a bit more integrated; they build webs together and attack prey in concert (Shear 1970, Buskirk 1981 [review], Nentwig 1985, Ward and Enders 1985), thereby handling larger prey. But the ultimate cooperative arthropods are the social hymenoptera, the bees, ants, and wasps.

Laymen usually consider members of these insect colonies subunits of a large organism centered on a queen. This is not entirely accurate. Workers in an ant colony are specialized for their roles, to be sure, but they are not genetically identical to the queen nor to one another. The questions of why they work for the good of the colony and, in particular, why they forego reproduction are not trivial. Hamilton (1963, reviewed in Wilson 1971) provided a convincing explanation of this, which is now universally accepted: by foregoing reproduction and adopting a specialized role, a worker can actually pass on more copies of its genes to the next generation, through the agency of its mother. In the hymenoptera, a female bears a greater relationship to her sisters than to her own offspring. This Darwinian explanation of insect sociality suggests that under certain circumstances a worker should reject its selfless role, for example when the queen begins to behave improperly, and this turns out to be true (Wilson 1971). When the queen, and members of other castes as well, deviate from normal behavior (a queen might cease

to produce male offspring, for example), normally passive workers turn on them and kill them. The illusion that these are mechanical, selfless societies is just that. The interesting question then emerges as to how these colonies are organized and how their members coordinate their activities. The richness of variety in the social structures of these insects is beyond the scope of this review (see Wilson 1971), but a few major points about cooperation emerge from the vast literature on the subject. The first is the question of caste differentiation.

The amount of differentiation ranges from none in some polistine wasps to four or five morphologically distinct castes in certain ants. It is not difficult to find ant colonies with a queen, drones, specialized soldiers, and both large and small (major and minor) workers. Within these castes there may be further role division based upon age. Oster and Wilson (1978) examined the evolutionary benefits of caste division and discovered two. The first is that the behavior of an individual colony member is dramatically simplified by comparison to an analogous solitary animal, allowing for morphological as well as behavioral specialization. This division of labor is inevitably accompanied by a loss in the total behavioral repertory of an individual, but the increase in the effectiveness with which the remaining behaviors are executed apparently compensates for it. The second key advance resulting from caste differentiation is the ability to conduct operations concurrently instead of sequentially. Responses to environmental change can be more prompt, thorough, and massive.

In arriving at their conclusions, Oster and Wilson relied heavily upon a key concept from the theory of reliability (Barlow and Proschan 1975), that a necessary job is more likely to be accomplished by a set of parallel devices than a set of serial devices (see their Figure 1.4, Appendix II). This implies that in a large group whose proper functioning requires the performance of a well-defined set of tasks, the evolution of role specialization is almost inevitable; moreover, the optimum number of castes is equal to the number of distinct tasks. This is an interesting concept that merits further thought. In particular, I wonder how large a group must be before role specialization becomes superior to role flexibility. And I am also concerned with the robustness of the assumption that one component is as good as many in permitting effective group functioning. This may be true of queens but is clearly untrue of workers. A colony with 5,000 workers will out-perform one with 5 or 500. Group performance is more than the mere absence of group failure.

How individual colony members operate is worth noting. Social insects apparently do not recognize one another as individuals. They recognize members of the colony and members of different castes, and they usually recognize different physiological categories within a caste.

But they do not recognize individuals. As a consequence there is no evidence that they form cliques or teams. Oster and Wilson assert that jobs are done less efficiently by redundant teams than by redundant components not organized into teams, and that this difference is reversed only if the level of coordination of team members is quite high, well beyond that of which social insects are capable. Be that as it may, social insects apparently do not exhibit highly integrated group behaviors. For that matter individual behaviors are not all that predictable. If a colony is building a nest, 100 individuals may be adding material while 20 remove it. The consequence is that the nest gets built, but speed and efficiency suffer somewhat. Wilson (1972) has proposed a stochastic theory of mass behavior to describe this phenomenon, in which each insect's behavior is captured as a series of transition probabilities between the different acts in its behavioral repertory. The whole colony's behavior is the statistical average of all these individual probability machines. This confers a degree of effective cooperation on the colony without any need for centralized control.

Probably the most commonly-cited example of cooperative behavior in animals is helping-at-the-nest. Brown's review of cooperation in animals (Brown 1983) dealt entirely with this phenomenon and nothing else. What happens is this: young animals, rather than attempt breeding, attach themselves to an adult breeding pair and participate in caring for the young of that pair. Usually the helpers are among the offspring of the breeding pair from previous years. The phenomenon has been observed in birds, mammals, and fish. Emlen (1984) and Gibbons (1987) provided reviews of the many species in which this behavior has been documented. Twenty years ago helpers were thought to emerge only from a pair's previous offspring and, therefore, to bear some genetic similarity to the new young they were helping. We now know of many cases where this is not true; kinship with the young is apparently not necessary for the evolution of helping. Two elements are common to all documented cases of helping, all such taxa show extended parental care and all have a surplus of breeding individuals. In other words it is difficult to rear young in these species, thus leading the parents to tolerance of helpers, and young adults have a hard time establishing themselves. Helping seems to be a strategy for positioning oneself for taking over a territory or mate if a dominant adult disappears.

The question of how helpers actually help seems to have been of relatively little interest, by comparison to the question of why they should help. As a consequence not much detailed information is available on their activities. Helpers may do little more than defend the territory against intruding conspecifics. In most cases they guard the nest against predators, freeing the parents for increased foraging activities. And in some cases helpers actively forage for the

young themselves. In any case their help is important. The improved reproductive success of breeding pairs with helpers is well documented (Brown 1983).

In terms of the degree of coordination, group hunting is probably the ultimate expression of cooperation in animals. Unfortunately, information on this class of behaviors is 99 parts anecdote to each part of good observational data. The phenomenon is so interesting to people, that more is written about it than is actually known. For example following a long wildlife tradition, Kleiman and Eisenberg (1973) asserted, with no evidence, that the underlying cause of sociality in the Carnivora was increased group effectiveness at killing large prey. This quickly became dogma. Packer (1986), in response, claimed that no evidence exists to demonstrate that lions are better hunters in groups than individually, that the real underlying evolutionary pressure for sociality in the species is group defense of territories and kills. Perhaps Packer is correct, but lack of evidence is certainly not evidence against something. At this time we simply have no idea whether lions are better hunters in groups.

What we do know is that in many cases of stalking in large carnivores, one or more group members have been observed to depart from a direct path to the prey, circle to the side, and wait in ambush. Presumably when the main body attacks, the prey will on occasion flee directly into the ambush. This has been observed in wolves (Mech 1970), lions (Schaller 1972), and coyotes (Robinson 1952, Hamlin 1979). In coyotes this behavior seems to be directed toward getting a female ungulate away from her calf or fawn. One coyote lies flat while a second moves around to the side or back of the female. The first coyote then rushes, drawing the counterattacking female away. The second coyote runs in and grabs the fawn. Whether this occurs often and whether coyotes rely much upon other style of group hunting is currently a source of argument (Bowen 1981; Beckoff and Wells 1986).

A great deal of speculation exists in the literature about cooperation during chasing. I can assert with some confidence that no reliable accounts exist on cooperation among social carnivores, once an attack proceeds to full pursuit.

A large number of studies of birds of prey claim cooperative hunting. Again I must discount all of these but one, but this one happens to be the single best-documented account of cooperative hunting in any species of vertebrate. Hector (1986) watched Aplomado falcons hunting in Mexico. He monitored 357 hunts by 18 different pairs of birds. Two falcons chased the same prey in 102 of those hunts. Perhaps more importantly, both falcons participated in 66% of the bird hunts. Sharp sex roles were observed in pursuit, and these were consistent with

body size differences. As is typical in falcons, females are larger (50% in this case) and less maneuverable. Attacks on birds were usually initiated by males. If the female was inattentive at that point, the male would give sharp "chip" or "cheep" vocalizations until she followed. When attacking flying prey, the two birds maintained horizontal parallel flight paths. When attacking prey in trees, females would fly close to the group and then ascend directly up into the tree while the male hovered overhead. The female pursued the prey actively through trees and on the ground while the male maintained a higher position, attacking whenever the quarry broke cover. When the prey were birds, solo hunts were 21% successful, and tandem hunts were 45% successful. Perhaps as surprising as these data is the fact that Hector attempted a general definition of cooperative hunting:

If hunts are cooperative ventures then they should show some of the following characteristics: (1) individuals tend to hunt together instead of hunting alone; (2) group members usually select the same prey animals for pursuit; (3) some division of labor occurs during hunts; (4) some signal (or signals) is used to coordinate movements of participants; (5) defendable food is shared among participants; and (6) individuals monitor each other's movements during hunts. (p. 248)

To my knowledge this is the first time anybody has ever published a definition of cooperative hunting, although the definition is a bit hazy. How many of these characteristics, for example, comprise "some"?

CONCLUSION

Several points emerge from this review. I shall recapitulate these and discuss their relevance to machine intelligence. The first is that the overwhelming emphasis of both human and animal behaviorists has been on the fact of cooperation rather than its nature. This is not entirely irrelevant to questions of autonomous machine cooperation. Components of a single machine can safely be designed to be completely altruistic; after all, if any one of them fails the entire mechanism breaks down. But as soon as the components become separate autonomous entities, the degree of interdependence becomes dramatically less. The consequence is that designers must begin to entertain questions of whether to cooperate or act to preserve oneself; it would not be intelligent for a robot to destroy itself for the sake of a very minor group gain. The question of how one balances group benefits against individual benefits has been treated in exceedingly fine detail in the literature and need not be reinvented.

The second major point I infer from the literature is that humans are complex creatures with occasionally byzantine goal structures and tactics. Their motives in social situations are rarely straightforward and obvious, even to themselves. Their cooperative tendencies are socially conditioned to a large extent and are remarkably fluid. While the fundamentally cooperative nature of human societies seems clear, the actual mechanisms by which cooperation operates are not. The use of human behavior as metaphor seems risky, perhaps even arrogantly so, at this point.

The third point is that the basic behavioral elements of cooperation, coordination, communication, and so forth, seem not to have been identified (Franklin and Harmon, 1987). To some extent this traces back to research traditions, such as the games approach, and to some extent it reflects simple negligence. I consider it unlikely that any group of behavioral components will be found to be common to all forms of cooperation, but we can certainly do better than Hector's preliminary attempt.

The fourth point is that cooperation seems to have gone along two quite different evolutionary tracks in animals. The track most familiar to us emphasizes individual recognition, role specialization, communication, and, to some extent, centralized control. This is reflected in vertebrate helping behaviors and cooperative hunting. The alternative track is that typified by

social insects, in which role specialization exists to a limited extent but no trace can be found of individual recognition, sophisticated communication, or centralized control. Neither form of cooperation is clearly ecologically dominant; both have strengths and weaknesses.

The insect strategy has the benefit of extraordinary resilience. No single individual, save in some species the queen, is indispensable. While a colony is impaired by the destruction of a majority of its members, it is seldom destroyed. Redundancy seems to be the key to this robustness. Each animal is inexpensive and quickly replaced.

The tasks these insects undertake are few, well-delineated, and hard-wired into the genes of each worker. As a consequence no need exists for centralized control of their operations. When workers sense the colony's need for food, they go out to forage. When workers sense a breach in the colony's surrounding structure, they begin repair behaviors. When workers sense a foreign insect in the colony, they attack it. The lack of coordination of workers in these tasks is more than compensated by the speed and massiveness of their responses.

Problems arise when a new task arises that is not in the colony's behavioral repertory. The biological response at this point is genetic and slow rather than behavioral and fast. And there seem to be performance limitations to this approach. While arm' ants are unstoppable predators, they do not catch healthy birds and mammals nor even alert mobile insects. This brings one to consideration of the vertebrate strategy.

Birds and mammals are the result of evolutionary pressures leading to larger, more expensive, and more capable bodies and to enhanced neural and sensory capacities. Cooperation is achieved by fewer individuals acting in a much more coordinated fashion. I suspect that the roles played in cooperation are largely learned rather than hard-wired. The degree of centralized coordination and control is unknown, although I suspect it is not high. Some role specialization occurs, although again we have no idea how this arises or if roles are at all flexible. If past research in animal behavior is any guide, further examination of this class of behaviors in vertebrates will almost certainly reveal subtleties of coordination heretofore unsuspected.

Which of the two types of cooperation, vertebrate or invertebrate, can be said to be more successful? This is a difficult question, although it seems to me that either provides a superior metaphor for machines when compared to humans. If one emphasizes reliability and resilience, the insect model is clearly superior. This implies the building of a cooperative group consisting of many inexpensive machines, each with limited capabilities. If one wants maximum

performance from an expensive machine, then the vertebrate model is obviously superior. In either case a good deal more observational work needs to be done on the fine-scale behavior of interest before any animal metaphor can be considered properly developed.

REFERENCES

A number of the references in this list are not mentioned in the review. They are included here because they make minor points that may be of interest. The conceptual value I attach to each paper is noted; papers of little or no conceptual value were excluded from the list.

Abric, J. C. 1982. Cognitive processes underlying cooperation: the theory of social representation. pp. 73-94. In Derlega, V. J. and Grzelak, J. L. (eds.) Cooperation and helping behaviors: theories and research. Academic Press, NY.

value = low

Barlow, R. E. and Proschan, F. 1967. Mathematical theory of reliability. Wiley, NY.

value = ?

Bartholdi, J. J. III, Butler, C. A., and Trick, M. A. 1986. More on the evolution of cooperation. J. Confl. Resol. 30:129-140.

value = low (another mathematical treatment of the prisoner's dilemma)

Bekoff, M. and Wells, M. C. 1986. Social ecology and behavior of coyotes. Adv. Study Behav. 16:251-338.

value = low

Ben-Yoav, O. and Pruitt, D. G. 1984. Resistance to yielding and the expectation of cooperative future interaction in negotiation. J. Exper. Soc. Psych. 20:323-335.

value = moderate

Bender, T. A. 1986. Monitoring and the transfer of individual problem solving. Cont. Educ. Psych. 11:161-169.

value = low

Bethlehem, D. W. 1982. Anthropological and cross-cultural perspectives. pp. 250-268. In Colman, A. M. (ed.) Cooperation and competition in humans and animals. Van Nostrand Rheinhold, Berks.

value = moderate-high

Bowen, W. D. 1981. Variation in coyote social organization: the influence of prey size. *Can. J. Zool.* 59:639-652.

value = low

Brady, J. E., Newcomb, A. F. and Hartup, W. W. 1983. Context and companion's behavior as determinants of cooperation and competition in school-age children. *J. Exp. Child Psych.* 36:396-412.

value = low

Braver, S. L. and Wilson, L. A., II. 1986. Choices in social dilemmas: effects of communication within subgroups. *J. Confl. Resol.* 30:51-62.

value = low

Brown, J. L. 1983. Cooperation: a biologist's dilemma. *Adv. Study Behav.* 13:1-39.

value = moderate

Buskirk, R. 1981. Sociality in arachnida. pp. 282-369. In Herman, H. (ed.) *Social insects*. Academic Press, NY.

value = low-moderate

Caraco, T. and Wolf, L. L. 1975. Ecological determinants of group sizes of foraging lions. *Amer. Natur.* 109:343-352.

value = low

Chelazzi, G., Deneuborg, J. L., and Focardi, S. 1985. Cooperative interactions and environmental control in the intertidal clustering of *Nerita textile*(Gastropoda: prosobranchia). *Behav.* 90:151-166.

value = low

Clark, C. W. 1987. The lazy, adaptable lions: a Markovian model of group foraging. *Anim. Behav.* 35:361-368.

value = low (game theory analysis of group size)

Collopy, M. W. 1983. Foraging behavior and success of golden eagles. *Auk* 100:747-749.

value = low (an account of a "group" attack on a duck)

Colman, A. M. 1982. Experimental games. pp. 113-140. In, Colman, A. M. (ed.) *Cooperation and competition in humans and animals*. Van Nostrand Rheinhold, Berks.

value = high

Colman, A. M. 1982. Theory of voting. pp. 89-112. In, Colman, A. M. (ed.) *Cooperation and competition in humans and animals*. Van Nostrand Rheinhold, Berks.

value = low

Cook, H. and Sloane, J. 1985. Locus of control and cooperative behavior in 10-year-old children. *J. Soc. Psych.* 125:619-630.

value = low

Cosden, M., Pearl, R. and Bryan, T. H. 1985. The effects of cooperative and individual goal structures on learning disabled and nondisabled students. *Excep. Childr.* 52:103-114.

value = low

Daniel, W. J. 1942. Cooperative problem-solving in rats. *J. Comp. Psych.* 34:361-368.

value = low

Daniel, W. J. 1947. How cooperatively do individual rats solve a problem in a social situation? *Amer. Psych.* 2:316.

value = low

Derlega, V. J. and Grzelak, J. L. 1982. Cooperation and helping behaviors: theories and research. Academic Press, NY.

value = high

Deutsch, M. 1982. Interdependence and psychological orientation. pp. 15-42. In Derlega, V. J. and Grzelak, J. L. (eds.) Cooperation and helping behaviors: theories and research. Academic Press, NY.

value = low

Emlen, S. T. 1984. Cooperative breeding in birds and mammals. pp. 305-339. In. Krebs, J. R. and Davies, N. B. (eds.) Behavioural ecology: an evolutionary approach. 2nd ed. Sinauer Assoc.

value = moderate

Espinosa, J. A. and Garza, R. T. 1985. Social group salience and interethnic cooperation. *J. Exp. Soc. Psych.* 21:380-392.

value = low

Ewer, R. F. 1973. The carnivores. Cornell Univ. Press, NY.

value = low

Franklin, R. F. and Harmon, L. A. 1987. Elements of cooperative behavior. IR&D Final Report, ERIM, Ann Arbor, Mi.

value = high

Gardner, R. M., Corbin, T. L. and Beltramo, J. S. 1984. The prisoner's dilemma game and cooperation in the rat. *Psych. Rep.* 55:687-696.

value = low

Georgas, J. 1985. Group interactions and problem solving under cooperative, competitive, and individual conditions with Greek children. *Genet. Soc. Gen. Psych. Mono.* 111:351-361.

value = low

Georgas, J. 1985. Cooperative, competitive and individual problem-solving in sixth grade Greek children. *Eur. J. Soc. Psych.* 15:67-77.

value = low-moderate

Ghent, A. W. 1960. A study of the group-feeding behaviour of larvae of the Jack Pine Sawfly, *Neodiprion pratti banksianae*. *Roh. Behav.* 16:110-148.

value = low

Ghiselin, M. 1974. The economy of nature and the evolution of sex. Univ. Cal. Press, Berkeley, Cal.

value = low

Gibbons, D. W. 1987. Juvenile helping in the moorhen. *Gallinula chloropus*. *Anim. Behav.* 35:170-181.

value = low

Gray, B. 1985. Conditions facilitating interorganizational collaboration. *Human Relat.* 38:911-936.

value = low

Grzelak, J. L. 1982. Preference and cognitive processes in interdependence situations: a theoretical analysis of cooperation. pp. 95-122. In. Derlega, V. J. and Grzelak, J. L. (eds.) *Cooperation and helping behavior: theories and research*. Academic Press, NY.

value = moderate

Hamilton, W. D. 1963. The evolution of altruistic behavior. *Amer. Natur.* 97:354-356.

value = low

Hamlin, K. L. 1979. Cooperation by coyote pairs attacking mule deer fawns. *J. Mammal.* 60:849-850.

value = low

Hector, D. P. 1986. Cooperative hunting and its relationship to foraging success and prey size in an avian predator. *Ethol.* 73:247-257.

value = high

Heretick, D. M. L. 1984. Trust-suspicion and gender differences in interpersonal functioning. *J. Res. Person.* 18:27-40.

value = low

Hughes, A. L. 1984. Some methods for analyzing the structure and behavior of human kin groups. *Eth. Sociobiol.* 5:179-192.

value = low

Janis, I. L. 1982. *Groupthink: psychological studies of policy decisions and fiascoes*. Houghton Mifflin, Boston.

value = low-moderate

Johnson, R. T., Johnson, D. W. and Stanne, M. B. 1985. Effects of cooperative, competitive, and individualistic goal structures on computer-assisted instruction. *J. Ed. Psych.* 77:668-677.

value = low

Kleiman, D. G. and Eisenberg, J. F. 1973. Comparisons of canid and felid social systems from an evolutionary perspective. *Anim. Behav.* 21:637-659.

value = low

Latane', B., Nida, S. A. and Wilson, D. W. 1981. The effects of group size on helping behavior. pp. 287-314. In Rushton, J. P. and Sorrentino, R. M. (eds.) *Altruism and helping behavior*. Erlbaum Assoc., NJ.

value = moderate

Laughlin, P. R. and Futoran, G. C. 1985. Collective induction: social combination and sequential transition. *J. Person. and Soc. Psych.* 48:608-613.

value = low

Lawrence, B. M. 1984. Conversation and cooperation: child linguistic maturity, parental speech, and helping behavior. *Child Dev.* 55:1926-1935.

value = low-moderate

Liebrand, W. B. G. 1984. The effect of social motives, communication and group size on behaviour in an N-person multi-stage mixed-motive game. *Eur. J. Soc. Psych.* 14:239-264.

value = low

Lindskold, S., Betz, B. and Walters, P. S. 1986. Transforming competitive or cooperative climates. *J. Confl. Resol.* 30:99-114.

value = low

Lynn, M. and Oldenquist, A. 1986. Egoistic and nonegoistic motives in social dilemmas. Amer. Psych. 41:529-534.

value = low

MacDonald, D. W. 1983. The ecology of carnivore social behavior. Nature. 301:379-384.

value = low

Malcolm, J. R. and Marten, K. 1982. Natural selection and the communal rearing of pups in African wild dogs (*Lycaon pictus*). Behav. Ecol. Sociobiol. 10:1-13.

value = low

McClintock, C. G. and Van Avermaet, E. 1982. Social values and rules of fairness: a theoretical perspective. pp. 43-71. In Derlega, V. J. and Grzelak, J. L. (eds.) Cooperation and helping behaviors: theories and research. Academic Press, NY.

value = moderate

McCleod, M. and Nowicki, S., Jr. 1985. Cooperative behavior as a function of interpersonal style in preschoolers. J. Person. 55:36-

value = low

Mech, D. L. 1970. The wolf: the ecology and behavior of an endangered species. Univ. Minn. Press, Mpls., Mn.

value = low

Nentwig, W. 1985. Social spiders catch larger prey: a study of *Anelosimus eximius* (Aranaea: Theridiidae). Behav. Ecol. Sociobiol. 17:79-85.

value = low

Oster, G. F. and Wilson, E. O. 1978. Caste and ecology in the social insects. Princeton Univ. Press, Princeton, NJ.

value = high

Packer, C. 1986. The ecology of sociality in felids. pp. 429-451. In Rubenstein, D. I. and Wrangham, R. W. (eds.) Ecological aspects of social evolution. Princeton Univ. Press, Princeton, NJ.

value = low-moderate

Pigott, H. E. and Heggie, D. L. 1985. Interpreting the conflicting results of individual versus group contingencies in classrooms: the targeted behavior as a mediating variable. Child. Fam. Behav. Ther. 7:1.

value = low

Rabbie, J. M. 1982. The effects of intergroup competitions and cooperation on intragroup and intergroup relationships. pp. 123-149. In Derlega, V. J. and Grzelak, J. L. (eds.) Cooperation and helping behaviors: theories and research. Academic Press, NY.

value = low-moderate

Robinson, W. B. 1952. Some observations on coyote predation in Yellowstone National Park. *J. Mammal.* 33:470-476.

value = low

Schaller, G. B. 1972. The Serengeti lion: a study of predator-prey relationships. Univ. Chicago Press, Chicago, Ill.

value = low

Schellenberg, J. A. 1985. Behavioral differences between structures with rank-order equivalence. *J. Confl. Resol.* 29:531-543.

value = low

Schmitt, D. R. 1984. Interpersonal relations: cooperation and competition. *J. Exp. An. Behav.* 42:377-383.

value = moderate-high

Schumacher, J. A. 1984. The model of a human being in public choice theory. *Amer. Behav. Sci.* 28:211-231.

value = moderate

Shear, W. A. 1970. The evolution of social phenomena in spiders. *Bull. Brit. Arachn. Soc.* 1:65-76.

value = low

Smith, E. A. 1985. Inuit foraging groups: some simple models incorporating conflicts of interest, relatedness, and central-place sharing. *Eth. Sociobiol.* 6:27-47.

value = low-moderate

Stingle, S. F. and Cook, H. 1985. Age and sex differences in the cooperative and noncooperative behavior of pairs of American children. *J. Psych.* 119:335-345.

value = low

Stockdale, D. F., Galejs, I. and Wolins, L. 1983. Cooperative-competitive preferences and behavioral correlates as a function of sex and age of school-age children. *Psych. Rep.* 53:739-750.

value = low

Stults, D. M. and Meese, L. A. 1985. Behavioral consistency: the impact of public versus private statements of intentions. *J. Soc. Psych.* 125:277-278.

value = low

Surazka, W. 1986. Normative and instrumental functions of equity criteria in individual choices of the input-payoff distribution pattern. *J. Confl. Resol.* 30:532-550.

value = low

Taborovsky, M. 1985. Breeder-helper conflict in a cichlid fish with broodcare helpers: an experimental analysis. *Behav.* 95:45-75.

value = low

Turner, J. C., et al. 1984. Failure and defeat as determinants of group cohesiveness. *Brit. J. Soc. Psych.* 23:97-111.

value = low

Ward, P. I. and Enders, M. M. 1985. Conflict and cooperation in the group feeding of the social spider *Stegodyphus mimosarum*. *Behav.* 94:167-182.

value = low

Wilson, E. O. 1971. *The insect societies*. Harvard Univ. Press, Cambridge, Mass.

value = high

Yamagishi, T. 1986. The provision of a sanctioning system as a public good. *J. Person. Soc. Psych.* 51:110-116.

value = low

~~SERIM~~

DISTRIBUTION LIST

	<u>Copies</u>
Dr. William E. Isler DARPA 1400 Wilson Blvd. Arlington, VA 22209	2
Tibor F. Czako, AMSTA-RSS U.S. Army Tank-Automotive Command Warren, MI 48397-5000	2
Commander Defense Technical Information Center Bldg. 5, Cameron Station ATTN: DDAC Alexandria, VA 22304-9990	12
Manager Defense Logistics Studies Information Exchange ATTN: AMXMC-D Fort Lee, VA 23801-6044	2
Commander U.S. Army Tank-Automotive Command ATTN: AMSTA-DDL Warren, MI 48397-5000	2
Commander U.S. Army Tank-Automotive Command ATTN: AMSTA-CF (Mr. G. Orlicki) Warren, MI 48397-5000	1
Commander U.S. Army Tank-Automotive Command ATTN: AMSTA-TMS Warren, MI 48397-5000	1
Director AMSAA ATTN: AMXSY-MP (Mr. Cohen) Aberdeen Proving Ground, MD 21005-5071	1